arXiv:1702.04322v1 [cs.CC] 14 Feb 2017

Parameterized Algorithms for Recognizing Monopolar and 2-Subcolorable Graphs*

Iyad Kanj[1], Christian Komusiewicz[2], Manuel Sorge[3], and Erik Jan van Leeuwen[4]

[1]School of Computing, DePaul University, Chicago, USA, ikanj@cs.depaul.edu
[2]Institut für Informatik, Friedrich-Schiller-Universität Jena, Germany, christian.komusiewicz@uni-jena.de
[3]Institut für Softwaretechnik und Theoretische Informatik, TU Berlin, Germany, manuel.sorge@tu-berlin.de
[4]Max-Planck-Institut für Informatik, Saarland Informatics Campus, Saarbrücken, Germany, erikjan@mpi-inf.mpg.de

February 15, 2017

Abstract

A graph G is a (Π_A, Π_B)-graph if $V(G)$ can be bipartitioned into A and B such that $G[A]$ satisfies property Π_A and $G[B]$ satisfies property Π_B. The (Π_A, Π_B)-Recognition problem is to recognize whether a given graph is a (Π_A, Π_B)-graph. There are many (Π_A, Π_B)-Recognition problems, including the recognition problems for bipartite, split, and unipolar graphs. We present efficient algorithms for many cases of (Π_A, Π_B)-Recognition based on a technique which we dub inductive recognition. In particular, we give fixed-parameter algorithms for two NP-hard (Π_A, Π_B)-Recognition problems, Monopolar Recognition and 2-Subcoloring. We complement our algorithmic results with several hardness results for (Π_A, Π_B)-Recognition.

1 Introduction

A (Π_A, Π_B)-*graph*, for graph properties Π_A, Π_B, is a graph $G = (V, E)$ for which V admits a partition into two sets A, B such that $G[A]$ satisfies Π_A and $G[B]$ satisfies Π_B. There is an abundance of (Π_A, Π_B)-graph classes, and important ones include *bipartite graphs* (which admit a partition into two independent

*A preliminary version of this paper appeared in SWAT 2016, volume 53 of LIPIcs, pages 14:1–14:14. Iyad Kanj and Manuel Sorge gratefully acknowledge the support by Deutsche Forschungsgemeinschaft (DFG), project DAPA, NI 369/12. Christian Komusiewicz gratefully acknowledges the support by Deutsche Forschungsgemeinschaft (DFG), project MAGZ, KO 3669/4-1.

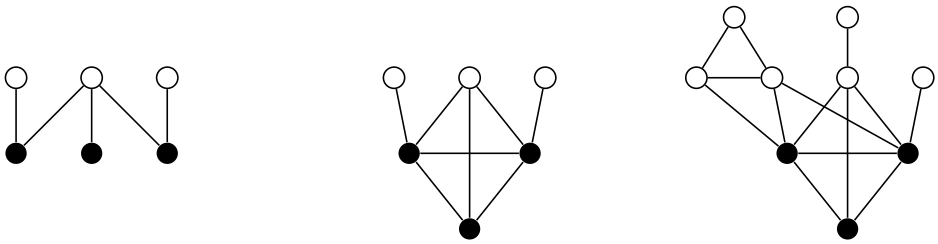

Figure 1: Three examples of (Π_A, Π_B)-graphs, where the coloring gives a (Π_A, Π_B)-partition. The vertices of A are black and the vertices of B are white. Left: in bipartite graphs A and B are independent sets. Center: in split graphs A is a clique and B is an independent set. Right: in unipolar graphs A is a clique and B induces a cluster graph.

sets), *split graphs* (which admit a bipartition into a clique and an independent set), and *unipolar graphs* (which admit a bipartition into a clique and a cluster graph). Here a *cluster graph* is a disjoint union of cliques. An example for each of these classes is shown in Figure 1.

The problem of recognizing whether a given graph belongs to a particular class of (Π_A, Π_B)-graphs is called (Π_A, Π_B)-Recognition, and is known as a *vertex-partition* problem. Recognition problems for (Π_A, Π_B)-graphs are often NP-hard [1, 12, 19], but bipartite, split, and unipolar graphs can all be recognized in polynomial time [11, 15, 22, 23, 28]. With the aim of generalizing these polynomial-time algorithms, we study the complexity of recognizing certain classes of (Π_A, Π_B)-graphs, focusing on two particular classes that generalize split and unipolar graphs, respectively. To achieve our goals, we formalize a technique, which we dub inductive recognition, that can be viewed as an adaptation of iterative compression to recognition problems. We believe that the formalization of this technique will be useful in general for designing algorithms for recognition problems.

Inductive Recognition The *inductive recognition* technique, described formally in Section 3, can be applied to solve the (Π_A, Π_B)-Recognition problem for certain hereditary (Π_A, Π_B)-graph classes. Intuitively, the technique works as follows. Suppose that we are given a graph $G = (V, E)$ and we are to decide its membership of the (Π_A, Π_B)-graph class. We proceed in iterations and fix an arbitrary ordering of the vertices; in the following, let $n := |V|$ and $m := |E|$. We start with the empty graph G_0, which trivially belongs to the hereditary (Π_A, Π_B)-graph class. In iteration i, for $i = 1, \ldots, n$, we recognize whether the subgraph G_i of G induced by the first i vertices of V still belongs to the graph class, assuming that G_{i-1} belongs to the graph class.

Inductive recognition is essentially a variant of the iterative compression technique [25], tailored to recognition problems. The crucial difference, however, is that in iterative compression we can always add the ith vertex v_i to the

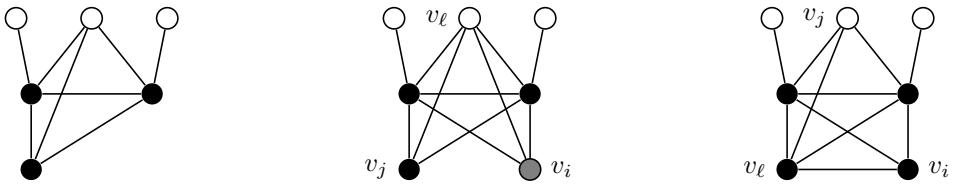

Figure 2: An example of the inductive step in inductive recognition. Left: a split graph G_{i-1} with a given partition into a clique A and an independent set B. Center: the partition for G_{i-1} cannot be directly extended to a partition for G_i since the vertex v_i has a nonneighbor v_j in A and a neighbor v_ℓ in B. Right: after deciding to put v_i in the clique A, we can repair the partition by moving v_j to the independent set B and v_ℓ to the clique A.

solution from the previous iteration to obtain a new solution (which we compress if it is too large). However, in the recognition problems under consideration, we cannot simply add v_i to one part of a two-partition (A, B) of G_{i-1}, where G_{i-1} is member of the graph class, and witness that G_i is still a member of the graph class: Adding v_i to A may violate property Π_A *and* adding v_i to B may violate property Π_B. An example for split graph recognition is presented in Figure 2. Here, we cannot add v_i to A or B to obtain another split graph, even if G_{i-1} is a split graph. Therefore, we cannot perform a 'compression step' as in iterative compression. Instead, we must attempt to add v_i to each of A and B, and then attempt to 'repair' the resulting partition in each of the two cases, by rearranging vertices, into a solution for G_i (if a solution exists). This idea is formalized in the inductive recognition framework in Section 3.

Monopolar Graphs and Mutually Exclusive Graph Properties The first (Π_A, Π_B)-Recognition problem that we consider is the problem of recognizing monopolar graphs, which are a superset of split graphs. A *monopolar graph* is a graph whose vertex set admits a bipartition into a cluster graph and an independent set; an example is shown in Figure 3. Monopolar graphs have applications in the analysis of protein-interaction networks [4]. The recognition problem of monopolar graphs can be formulated as follows:

> Monopolar Recognition
> **Input:** A graph $G = (V, E)$.
> **Question:** Does G have a *monopolar partition* (A, B), that is, can V be partitioned into sets A and B such that $G[A]$ is a cluster graph and $G[B]$ is an edgeless graph?

In contrast to the recognition problem of split graphs, which admits a linear-time algorithm [15], Monopolar Recognition is NP-hard. This motivates a parameterized complexity analysis of Monopolar Recognition. We consider the parameterized version of Monopolar Recognition, where the parameter

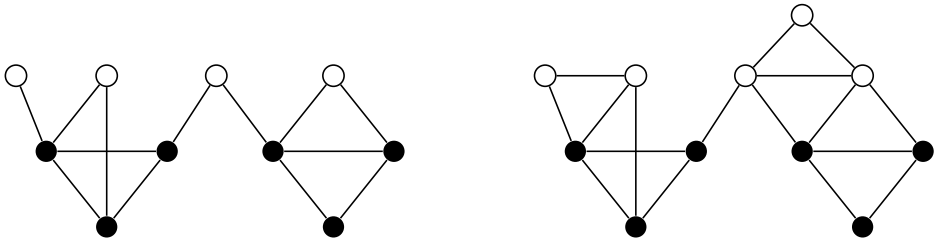

Figure 3: A monopolar and a 2-subcolorable graph with colorings that give a (Π_A, Π_B)-partition. The vertices of A are black and the vertices of B are white. Left: a monopolar graph where A induces a cluster graph and B is an independent set. Right: a 2-subcolorable graph where A and B induce cluster graphs.

k is an upper bound on the number of clusters in $G[A]$, and use inductive recognition to show the following:

Theorem 1.1. In $\mathcal{O}(2^k \cdot k^3 \cdot (n+m))$ time, we can decide whether G admits a monopolar partition (A, B) such that $G[A]$ is a cluster graph with at most k clusters.

The above result is a generalization of the linear-time algorithm for recognizing split graphs [15], which we can obtain by setting $k = 1$.

We observe that inductive recognition, and the ideas used to obtain the result in Theorem 1.1, can be generalized to (Π_A, Π_B)-Recognition problems in which properties Π_A and Π_B satisfy certain conditions of *mutual exclusivity*. To understand this notion of mutual exclusivity, consider the situation for split graphs. Intuitively, recognizing split graphs is easy for the following reason: If we consider any bipartition of a split graph G into a clique K and an independent set I, then this bipartition differs only very slightly from any other such bipartition of G. This is because at most one vertex of K may be part of any independent set of G and similarly, at most one vertex of I may be part of any clique in G. Thus, the two graph properties exclude each other to a large extent. In the situation of monopolar graphs, the two properties Π_A and Π_B defining monopolar graphs, are not mutually exclusive: $G[A]$ and $G[B]$ could be an edgeless graph of arbitrary order. However, in the parameterized Monopolar Recognition problem, we restrict $G[A]$ to contain at most k clusters. This restriction makes the properties mutually exclusive, as $G[A]$ may not contain an edgeless graph on $k+1$ vertices anymore. Hence, any two monopolar partitions with k clusters in $G[A]$ again differ only slightly if the parameter k is small. We generalize this observation in the following definition:

Definition 1.1. Two graph properties Π_A and Π_B are called *mutually d-exclusive* if there is no graph of order at least d that fulfills Π_A and Π_B.

For a pair (Π_A, Π_B) of mutually d-exclusive graph properties, we use inductive recognition to obtain the following result:

Theorem 1.2. If Π_A and Π_B are hereditary and mutually d-exclusive, and membership of Π_A and Π_B can be decided in polynomial time, then (Π_A, Π_B)-Recognition can be solved in $n^{2d+\mathcal{O}(1)}$ time.

Although Theorem 1.2 is quite general, there are many natural cases of (Π_A, Π_B)-Recognition to which it does not apply. Moreover, the degree of the polynomial in the running time depends on d, that is, the corresponding algorithm is an XP algorithm for the parameter d. Thus, while Theorem 1.2 can be applied to solve the parameterized Monopolar Recognition problem, the running time obtained is not practical, even for moderate values of k. In contrast, the much more efficient algorithm alluded to in Theorem 1.1 is tailored for parameterized Monopolar Recognition. Such an improvement from the XP algorithm implied by Theorem 1.2 to an FPT algorithm is sometimes unlikely, as we show in Section 4.

2-Subcolorings Next, we study *2-subcolorable graphs*; these are graphs for which the vertex set admits a bipartition into two cluster graphs [3], and thus are a superset of unipolar graphs; an example is shown in Figure 3. The recognition problem of 2-subcolorable graphs can be formulated as follows:

> 2-Subcoloring
> **Input:** A graph $G = (V, E)$.
> **Question:** Does G have a *2-subcoloring* (A, B), that is, can V be partitioned into sets A and B such that each of $G[A]$ and $G[B]$ is a cluster graph?

In contrast to the recognition problem of unipolar graphs, which admits a polynomial-time algorithm [11, 22, 28], 2-Subcoloring is NP-hard [1]. We consider 2-Subcoloring parameterized by the number of clusters in $G[A]$, and use inductive recognition to show the following:

Theorem 1.3. In $\mathcal{O}(k^{2k+1} \cdot nm)$ time, we can decide whether G admits a 2-subcoloring (A, B) such that $G[A]$ is a cluster graph with at most k clusters.

The above result can be seen as a generalization of the polynomial-time algorithms for recognizing unipolar graphs [11, 22, 28], which we can obtain by setting $k = 1$. We remark that one faces various technical difficulties when designing algorithms for parameterized 2-Subcoloring as it does not seem to yield to standard approaches in parameterized algorithms. This is a testament to the power of the inductive recognition technique, which adds to the arsenal of existing techniques for designing parameterized algorithms. Observe also that Theorem 1.2 does not apply to parameterized 2-Subcoloring.

Further Results We also consider the 2-Subcoloring problem parameterized by a weaker parameter: the total number of clusters in $G[A]$ and $G[B]$. This parameterization makes the problem amenable to a branching strategy that branches on the placement of the endpoints of suitably-chosen edges and

nonedges of the graph. In this way, we create partial 2-subcolorings (A', B') where each vertex in $V \setminus (A' \cup B')$ is adjacent to the vertices of exactly two partial clusters, one in each of $G[A']$ and $G[B']$. Then we show that whether such a partial 2-subcoloring extends to an actual 2-subcoloring of G can be tested in polynomial time via a reduction to 2-CNF-Sat. We prove the following result:

Theorem 1.4. In $\mathcal{O}(4^k \cdot k^2 \cdot n^2)$ time, we can decide whether G admits a 2-subcoloring (A, B) such that $G[A]$ and $G[B]$ are cluster graphs with at most k clusters in total.

Finally, we consider the parameter consisting of the total number of vertices in $G[A]$. We observe that a straightforward branching strategy yields a generic fixed-parameter algorithm for many (Π_A, Π_B)-Recognition problems.

Proposition 1.1. Let Π_A and Π_B be two hereditary graph properties such that membership of Π_A can be decided in polynomial time and Π_B can be characterized by a finite set of forbidden induced subgraphs. Then we can decide in $2^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$ time whether V can be partitioned into sets A and B such that $G[A] \in \Pi_A$, $G[B] \in \Pi_B$, and $|A| \leq k$.

We complement our results by presenting in Section 9 several hardness results showing that significant improvements over the algorithms presented in this paper for Monopolar Recognition, 2-Subcoloring, and the (Π_A, Π_B)-Recognition problem in general, are unlikely.

Related Work As mentioned, split graphs and unipolar graphs can be recognized in linear time [15] and polynomial time [11, 22, 28], respectively. In contrast, if Π_A and Π_B can be characterized by a set of connected forbidden subgraphs and Π_A is not the set of all edgeless graphs, then (Π_A, Π_B)-Recognition is NP-hard [12]. This implies the NP-hardness of Monopolar Recognition and 2-Subcoloring. Up to the authors' knowledge, the parameterized complexity of Monopolar Recognition and 2-Subcoloring has not been studied before. The known algorithms for both problems are not parameterized, and assume that the input graph belongs to a structured graph class; see [5, 6, 10, 20] and [3, 13, 14, 26], respectively. Recently, Kolay and Panolan [18] considered the problem of deleting k vertices or edges to obtain an (r, ℓ)-graph. For integers r, ℓ, a graph $G = (V, E)$ is an (r, ℓ)-*graph* if V can be partitioned into r independent sets and ℓ cliques. For example, $(2, 0)$-graphs are precisely bipartite graphs and $(1, 1)$-graphs are precisely split graphs. Observe that $(1, \cdot)$-graphs are *not* necessarily monopolar, because monopolar graphs do not allow edges between the cliques (as $G[A]$ is a cluster graph), whereas such edges are allowed in $(1, \cdot)$-graphs. These differences lead to substantially different algorithmic techniques. For example, since Kolay and Panolan [18] consider the deletion problem, they can use iterative compression in their work. Moreover, they consider $r, \ell < 3$, whereas the number of cliques k may be arbitrarily large in our setting. Nevertheless, during the development of our algorithms we were inspired by some of their gubernatorial observations.

Organization of the Article After describing the necessary graph-theoretic notations and giving a brief background on parameterized complexity in Section 2, we introduce inductive recognition in Section 3. In Section 4, we show our most general tractability result based on inductive recognition, the XP algorithm for mutually exclusive graph properties (Theorem 1.2). In Section 5, we give the linear-time fixed-parameter algorithm for Monopolar Recognition parameterized by the number k of cliques (Theorem 1.1). In Section 6, we present more general graph classes such that we can obtain FPT algorithms for their corresponding recognition problems. In Section 7, we turn to 2-Subcoloring parameterized by the smaller number of cliques between the two parts, and give an FPT algorithm for this problem based on inductive recognition (Theorem 1.3). In Section 8, we present FPT algorithms that are not based on inductive recognition, for recognition problems parameterized by weaker parameters (Proposition 1.1 and Theorem 1.4). Our hardness results are presented in Section 9. In Section 10, we summarize our findings and point out future research directions.

2 Preliminaries

For $n \in \mathbb{N}$, we let $[n] := \{1, \ldots, n\}$ denote the integers from 1 through n. We follow standard graph-theoretic notation [8]. Let G be a graph. By $V(G)$ and $E(G)$ we denote the vertex set and the edge set of G, respectively. The *order* of a graph G is $|V(G)|$. Throughout this work, let $n := |V(G)|$ and $m := |E(G)|$. For $X \subseteq V(G)$, $G[X]$ denotes the *subgraph of G induced by X*. For a vertex $v \in G$, $N(v)$ and $N[v]$ denote the open neighborhood and the closed neighborhood of v, respectively. The *degree* of a vertex $v \in G$, denoted $\deg(v)$, is $|N(v)|$. We say that a vertex v is *adjacent to a subset $X \subseteq V(G)$ of vertices* if v is adjacent to at least one vertex in X. For $X \subseteq V(G)$, we define $N(X) := (\bigcup_{v \in X} N(v)) \setminus X$ and $N[X] := \bigcup_{v \in X} N[v]$, and for a family $\mathcal{X}$ of subsets $X \subseteq V(G)$, we define $N(\mathcal{X}) := (\bigcup_{X \in \mathcal{X}} N(X)) \setminus (\bigcup_{X \in \mathcal{X}} X)$ and $N[\mathcal{X}] := \bigcup_{X \in \mathcal{X}} N[X]$. If S is any set of vertices in G, we write $G - S$ for the subgraph of G obtained by deleting all the vertices in S and their incident edges. For a vertex $v \in V(G)$, we write $G - v$ for $G - \{v\}$.

By P_3 we denote the graph that is a (simple) path on 3 vertices. We repeatedly use the following well-known fact:

Fact 2.1. A graph is a cluster graph if and only if it does not contain P_3 as an induced subgraph.

Let r, s be positive integers. The *Ramsey number* $R(r, s)$ is the smallest integer such that every graph of order at least $R(r, s)$ either contains a clique of r vertices or an independent set of s vertices. Ramsey's theorem [24] states that, for any two positive integers r, s, the number $R(r, s)$ exists. The following upper bound on $R(r, s)$ is known: $R(r, s) \leq \binom{r+s-2}{r-1}$.

A *parameterized problem* is a set of instances of the form (x, k), where $x \in \Sigma^*$ for a finite alphabet set Σ, and $k \in \mathbb{N}$ is the *parameter*. A parameterized problem

Q is *fixed-parameter tractable* (FPT), if there exists an algorithm that on input (x, k) decides if (x, k) is a yes-instance of Q in time $f(k)n^{\mathcal{O}(1)}$, where f is a computable function independent of $n = |x|$; an algorithm with this running time is called FPT *algorithm*. A hierarchy of fixed-parameter intractability, *the* W-*hierarchy* $\bigcup_{t\geq 0} \mathrm{W[t]}$, was introduced based on the notion of FPT *reduction*, in which the 0th level W[0] is the class FPT. It is commonly believed that W[1] $\neq$ FPT. A parameterized problem Q is in the parameterized complexity class XP, if there exists an algorithm that on input (x, k) decides if (x, k) is a yes-instance of Q in time $\mathcal{O}(n^{f(k)})$, where f is a computable function independent of $n = |x|$. For more discussion on parameterized complexity, we refer to the literature [7, 9].

We make also use of the *Exponential Time Hypothesis* (ETH). The ETH was formulated by Impagliazzo *et al.* [16], and states that k-CNF-Sat (for any $k \geq 3$) cannot be solved in subexponential time $2^{o(n)}$, where n is the number of variables in the input formula. Therefore, there exists a constant $c > 0$ such that k-CNF-Sat cannot be solved in time $\mathcal{O}(2^{cn})$. ETH has become a standard hypothesis in complexity theory for proving tight running time bounds results.

3 Foundations of Inductive Recognition

In this section, we describe the foundations of the general technique that we use to solve (Π_A, Π_B)-Recognition problems. The technique works in a similar way to the iterative compression technique by Reed et al. [25]. Let $\mathcal{G}$ be an arbitrary hereditary graph class (that is, if $G \in \mathcal{G}$, then $G' \in \mathcal{G}$ for every induced subgraph G' of G). We call an algorithm $\mathcal{A}$ an *inductive recognizer* for $\mathcal{G}$ if given a graph $G = (V, E)$, a vertex $v \in V$ such that $G - v \in \mathcal{G}$, and a membership certificate for $G - v \in \mathcal{G}$, algorithm $\mathcal{A}$ correctly decides whether or not $G \in \mathcal{G}$, and gives a membership certificate if $G \in \mathcal{G}$ (for example, in the case of recognizing monopolar graphs, it may be a string encoding a monopolar partition).

Theorem 3.1. Let $\mathcal{G}$ be an arbitrary hereditary graph class. Given an inductive recognizer $\mathcal{A}$ for $\mathcal{G}$, we can recognize whether a given graph $G = (V, E)$ is a member of $\mathcal{G}$ in time $\mathcal{O}(n + m) + \sum_{i=1}^{n} T(i)$, where $T(i)$ is the worst-case running time of $\mathcal{A}$ on a graph of order at most i.

Proof. We first sort the vertices in nondecreasing order of their vertex-degree in $\mathcal{O}(n + m)$ time (using Bucket Sort for example) to obtain a list $v_1, \ldots, v_n$, where $\deg(v_i) \leq \deg(v_{i+1})$, for $i = 1, \ldots, n-1$. Let $G_0 = (\emptyset, \emptyset)$ be the empty graph and $G_i := G[\{v_1, v_2, \ldots, v_i\}]$, for $i = 1, \ldots, n$. Since $\mathcal{G}$ is hereditary, G_0 is a member of $\mathcal{G}$ and we can easily compute a membership certificate of G_0 in $\mathcal{G}$. Then, for $i = 1, \ldots, n$, in order, we run $\mathcal{A}$ on (G_i, v_i) passing to $\mathcal{A}$ the certificate of membership of G_{i-1} in $\mathcal{G}$, to decide whether or not G_i is a member of $\mathcal{G}$, and produce a membership certificate in case it is, but only if G_{i-1} is a member of $\mathcal{G}$. If $\mathcal{A}$ decides that G_i is not a member of $\mathcal{G}$ for some $i = 1, \ldots, n$, then we answer that G is not a member of $\mathcal{G}$; this is correct, because $\mathcal{G}$ is hereditary. Otherwise,

we answer that G is a member of $\mathcal{G}$; the correctness of this answer follows from the correctness of $\mathcal{A}$. The bound on the running time is straightforward. □

For the purpose of this paper, we consider *parameterized inductive recognizers* for (Π_A, Π_B)-graphs. In addition to G and v, these recognizers take a nonnegative integer k as input. The above general theorem can then be instantiated as follows.

Corollary 3.1. Let k be a nonnegative integer, and let Π_A and Π_B be two hereditary graph properties. Let $\mathcal{G}_k$ be a class of (Π_A, Π_B)-graphs with an additional hereditary property that depends on k. Given a parameterized inductive recognizer $\mathcal{A}$ for $\mathcal{G}_k$, we can recognize whether or not a given graph $G = (V, E)$ is a member of $\mathcal{G}_k$ in time $\mathcal{O}(n + m) + \sum_{i=1}^{n} T(i, k)$, where $T(i, k)$ is the worst-case running time of $\mathcal{A}$ with parameter k on a graph of order at most i.

Corollary 3.2. Let k be a nonnegative integer and let Π_A and Π_B be two hereditary graph properties. Let $\mathcal{G}_k$ be a class of (Π_A, Π_B)-graphs with an additional hereditary property that depends on k. Given a parameterized inductive recognizer $\mathcal{A}$ for $\mathcal{G}_k$ that runs in time $f(k) \cdot \Delta$, where Δ is the maximum degree of the input graph and f is an arbitrary computable function, we can recognize whether or not a given graph $G = (V, E)$ is a member of $\mathcal{G}_k$ in time $f(k) \cdot (n + m)$.

Proof of Corollary 3.2. The proof of Theorem 3.1 implies that if the running time of $\mathcal{A}$ depends linearly on the maximum degree of the input graph and arbitrarily on k, then the total running time is linear for every fixed k, because $\deg(v_i)$ is an upper bound on the maximum degree of G_i. □

4 A General Application of Inductive Recognition

Recall that (Π_A, Π_B)-Recognition is NP-hard if Π_A and Π_B can be characterized by a set of connected forbidden induced subgraphs and, additionally, Π_A is not the set of all edgeless graphs [12]. While being quite general, this hardness result is not exhaustive and ideally, we would like to obtain a complexity dichotomy for (Π_A, Π_B)-Recognition. As a first application of inductive recognition and a step towards such a complexity dichotomy, we show how inductive recognition can be used to solve (Π_A, Π_B)-Recognition for hereditary mutually d-exclusive graph properties Π_A, Π_B, as defined in Definition 1.1 (Section 1).

To apply inductive recognition, we need to describe an inductive recognizer for (Π_A, Π_B)-graphs, that is, we need to give an algorithm for the following problem.

Inductive (Π_A, Π_B)-Recognition
Input: A graph $G = (V, E)$, a vertex $v \in V$, and a partition (A', B')

of $G' = G - v$ such that $G[A'] \in \Pi_A$ and $G[B'] \in \Pi_B$.
Question: Does V have a (Π_A, Π_B)-partition (A, B), that is, a partition such that $G[A] \in \Pi_A$ and $G[B] \in \Pi_B$?

For mutually d-exclusive graph properties, we can solve this problem by starting a search from the partition (A', B') of $G - v$. Herein, we can use the fact that the number of vertices that can be moved from A' to B and from B' to A are each upper-bounded by d, because $G[A'] \in \Pi_A$ and $G[B'] \in \Pi_B$. This implies that the partition (A, B) is determined to a large extent by the partition (A', B').

Lemma 4.1. If Π_A and Π_B are hereditary and mutually d-exclusive graph properties, and membership of Π_A and Π_B can be decided in polynomial time, then Inductive (Π_A, Π_B)-Recognition can be solved in time $n^{2d+\mathcal{O}(1)}$.

Proof. Assume there is a (Π_A, Π_B)-partition (A, B) of V. Since Π_A and Π_B are mutually d-exclusive, at most $d-1$ vertices of A' are contained in B and at most $d-1$ vertices of B' are contained in A. Consequently, we can decide whether G is a (Π_A, Π_B)-graph by the following algorithm. Consider each pair of $\tilde{A} \subseteq A'$ and $\tilde{B} \subseteq B'$ such that $|\tilde{A}| < d$ and $|\tilde{B}| < d$. Determine whether

$$((A' \cup \{v\} \cup \tilde{B}) \setminus \tilde{A}, (B' \cup \tilde{A}) \setminus \tilde{B})$$

is a (Π_A, Π_B)-partition and output it if this is the case. Otherwise, check whether

$$((A' \cup \tilde{B}) \setminus \tilde{A}, (B \cup \{v\} \cup \tilde{A}) \setminus \tilde{B})$$

is a (Π_A, Π_B)-partition and output it if this is the case. If both tests fail for all pairs of $\tilde{A}$ and $\tilde{B}$, then output that G is not a (Π_A, Π_B)-graph.

The correctness follows from the fact that the algorithm considers both possibilities of placing v in A or B, and all possibilities for moving vertices from A' to B and from B' to A. The running time is $n^{2d-2} \cdot n^{\mathcal{O}(1)} = n^{2d+\mathcal{O}(1)}$, since we consider altogether at most $2 \cdot n^{2d-2}$ different two-partitions, and for each two-partition the membership of the two parts in Π_A and Π_B can be determined in polynomial time. □

By combining Theorem 3.1 and Lemma 4.1, we immediately obtain the following.

Theorem 1.2. *If Π_A and Π_B are hereditary and mutually d-exclusive, and membership of Π_A and Π_B can be decided in polynomial time, then* (Π_A, Π_B)-Recognition *can be solved in $n^{2d+\mathcal{O}(1)}$ time.*

Two hereditary graph properties Π_A, Π_B are mutually d-exclusive for some integer d if and only if Π_A excludes some edgeless graph and Π_B excludes some clique, or vice versa: Clearly, the "only if"-part holds. For the "if"-part, we obtain the following upper bounds on d.

Proposition 4.1. Let Π_A and Π_B be hereditary graph properties. If Π_A excludes an edgeless graph of order s_a and Π_B excludes a complete graph of order s_b, then Π_A and Π_B are mutually $R(s_a, s_b)$-exclusive.

Proof. By the definition of Ramsey numbers, every graph of order $R(s_a, s_b)$ contains either an edgeless subgraph of order at least s_a, or a complete subgraph of order s_b. Thus, every graph of order at least $R(s_a, s_b)$ fulfilling Π_A contains a complete graph of order s_B and thus does not fulfill Π_B. Since Ramsey numbers are symmetric, every graph of order $R(s_a, s_b)$ fulfilling Π_B contains an edgeless subgraph of order s_a and thus does not fulfill Π_A. □

If Π_A and Π_B fulfill the conditions of the above proposition and are recognizable in polynomial time, then Theorem 1.2 applies.

Corollary 4.1. Let Π_A and Π_B be hereditary graph properties such that membership of Π_A and Π_B can be decided in polynomial time. If Π_A excludes a fixed edgeless graph and Π_B excludes a fixed complete graph, then (Π_A, Π_B)-Recognition can be solved in polynomial time.

Observe that if Π_A and Π_B both contain arbitrarily large edgeless graphs, then Π_A and Π_B are not mutually exclusive. Similarly, Π_A and Π_B are not mutually exclusive if both contain arbitrarily large complete graphs. As a consequence, Corollary 4.1 summarizes the applications of Theorem 1.2. A natural question is whether in Theorem 1.2 the dependency of the running time on d can be improved. A substantial improvement to $f(d_A, d_B) \cdot n^{\mathcal{O}(1)}$, however, is unlikely as we show in the remainder of this section.

A Note on Vertex Deletion Problems Theorem 1.2 has some applications for vertex deletion problems in undirected graphs which we illustrate with an example below. On the negative side, we find a graph property Π_A and a family $\mathcal{F}$ of graph properties such that Π_A and each $\Pi_B \in \mathcal{F}$ are mutually $d(\Pi_B)$-exclusive and (Π_A, Π_B)-Recognition is W[1]-hard with respect to $d(\Pi_B)$ if we consider Π_B as part of the input.

Consider the Vertex Cover problem where we are given a graph G and want to determine whether G has a vertex cover S of size at most k, that is, whether at most k vertices of G can be deleted such that the remaining graph is edgeless. This problem can be phrased as a (Π_A, Π_B)-Recognition problem: Π_A is the class of edgeless graphs and Π_B is the class of graphs of order at most k.

Let us, instead of using the standard parameter for Vertex Cover, the solution size k, parameterized by the smaller parameter ℓ, the "size of the maximum independent set over all solutions S". That is, we modify the problem by adding an additional integer ℓ to the input and we want to decide whether there is a solution of size at most k whose maximum independent set is of size ℓ. Call this problem Dense Vertex Cover. Clearly, ℓ can be arbitrarily smaller than k. A given instance (G, k, ℓ) of Dense Vertex Cover can again be formulated in terms of (Π_A, Π_B)-Recognition: As before, Π_A is the set of edgeless graphs, and Π_B is now the class of graphs of order at most k which have no independent set of size $\ell + 1$. Since Π_A excludes the complete graph on two vertices, Π_A and Π_B are mutually $(\ell+1)$-exclusive. Hence, Theorem 1.2 implies

for every fixed ℓ a polynomial-time algorithm for Dense Vertex Cover, in other words, an XP algorithm for the parameter ℓ.

Ideally, we would like to replace this XP algorithm by an FPT algorithm. This, however, is unlikely, as the following proposition shows.

Proposition 4.2. Dense Vertex Cover parameterized by ℓ is W[1]-hard.

Proof sketch. Consider the Partitioned Independent Set problem where we are given a graph $G = (V, E)$ with a vertex partition $V_1, \ldots, V_t$ such that each part V_i induces a clique in G and the task is to decide whether G has an independent set of size t. Equivalently, we may ask for a vertex cover of size $n-t$. Partitioned Independent Set is well-known to be W[1]-hard parameterized by the size t of the desired independent set. Furthermore, the parameter ℓ is upper-bounded by t in Partitioned Independent Set instances because the maximum independent set size in G is at most t. □

We have noted above that each instance of Dense Vertex Cover can be solved by a call to the algorithm in Theorem 1.2 for two mutually $(\ell + 1)$-exclusive graph properties. Combining this with Proposition 4.2, we obtain the following corollary.

Corollary 4.2. Unless FPT = W[1], the running time in Theorem 1.2 cannot be improved to $f(d) \cdot n^{\mathcal{O}(1)}$.

Nevertheless, Theorem 1.2 and Corollary 4.1 imply XP algorithms for many vertex deletion problems when the parameter is the size of the independent set of the solution; examples of such applications are Feedback Vertex Set (where the remaining graph is a forest), Bounded Degree Deletion (where the remaining graph has degree at most r for some constant r), and Planar Vertex Deletion.

5 An FPT Algorithm for Monopolar Recognition

Monopolar Recognition is the special case of (Π_A, Π_B)-Recognition where Π_A is the set of cluster graphs and Π_B is the set of edgeless graphs. Here, we consider Monopolar Recognition with the number k of clusters as a parameter. This further restricts Π_A: by bounding the parameter k we constrain Π_A to be the set of cluster graphs with at most k clusters. Thus, the graphs in Π_A cannot contain an edgeless graph of order $k + 1$ as subgraph. In other words, every graph in Π_A that has order at least $k + 1$ contains at least one edge. Altogether this implies the following.

Fact 5.1. If Π_A is the set of cluster graphs with at most k clusters and Π_B is the set of edgeless graphs, then Π_A and Π_B are mutually $(k + 1)$-exclusive.

Thus, Π_A and Π_B fulfill the conditions of Theorem 1.2 for each fixed k which implies an XP algorithm for Monopolar Recognition parameterized by k. In this section, we give a linear-time FPT algorithm for Monopolar Recognition parameterized by the number of clusters k.

Throughout, given a graph $G = (V, E)$ and a nonnegative integer k, we say that a monopolar partition (A, B) of G is *valid* if $G[A]$ is a cluster graph with at most k clusters. Using Corollary 3.2, it suffices to give a parameterized inductive recognizer for graphs with a valid monopolar partition. That is, we need to solve the following problem in time $f(k)\cdot\Delta$, where f is some computable function and Δ the maximum degree of G:

> Inductive Monopolar Recognition
> **Input:** A graph $G = (V, E)$, a vertex $v \in V$, and a valid monopolar partition (A', B') of $G' = G - v$.
> **Question:** Does G have a valid monopolar partition (A, B)?

In the following, we fix an instance of Inductive Monopolar Recognition with a graph $G = (V, E)$, a vertex $v \in V$, and a valid monopolar partition (A', B') of $G' = G - v$.

To find a valid monopolar partition (A, B) of G, we try the two possibilities of placing v in A or placing v in B. More precisely, in the first case, we start a search from the bipartition $(A' \cup \{v\}, B')$, and in the second case, we start a search from the bipartition $(A', B' \cup \{v\})$. Neither of these two partitions is necessarily a valid monopolar partition of G. The search strategy is to try to "repair" a candidate partition by moving few vertices from one part of the partition to the other part. During this process, if a vertex is moved from one part to the other, then it will never be moved back. To formalize this approach, we introduce the notion of constraints.

Definition 5.1. A *constraint* $\mathcal{C} = (A^{\mathcal{C}}_*, A^{\mathcal{C}}_P, B^{\mathcal{C}}_*, B^{\mathcal{C}}_P)$ is a four-partition of V such that $A^{\mathcal{C}}_* \subseteq A'$ and $B^{\mathcal{C}}_* \subseteq B'$. The vertices in $A^{\mathcal{C}}_P$ and $B^{\mathcal{C}}_P$ are called *permanent* vertices of the constraint. A constraint $\mathcal{C} = (A^{\mathcal{C}}_*, A^{\mathcal{C}}_P, B^{\mathcal{C}}_*, B^{\mathcal{C}}_P)$ is *fulfilled* by a vertex bipartition (A, B) of G if (A, B) is a valid monopolar partition of G such that $A^{\mathcal{C}}_P \subseteq A$ and $B^{\mathcal{C}}_P \subseteq B$.

The permanent vertices in $A^{\mathcal{C}}_P$ and $B^{\mathcal{C}}_P$ in the above definition will correspond to those vertices that were moved during the search from one part to the other part. The following observation is straightforward:

Fact 5.2. Each valid monopolar partition (A, B) of G fulfills either $(A', \{v\}, B', \emptyset)$ or $(A', \emptyset, B', \{v\})$.

We call the two constraints in Fact 5.2 the *initial constraints* of the search. We solve Inductive Monopolar Recognition by giving a search-tree algorithm that determines for each of the two initial constraints whether there is a partition fulfilling it. The root of the search tree is a dummy node that has two children, associated with the two initial constraints. Each non-root node in the search tree is associated with a constraint $\mathcal{C}$, and the algorithm searches

for a solution that fulfills $\mathcal{C}$. To this end, the algorithm applies reduction and branching rules that find vertices that in *every* valid monopolar partition (A, B) fulfilling $\mathcal{C}$ are in $A^{\mathcal{C}}_* \cap B$ or $B^{\mathcal{C}}_* \cap A$; that is, these vertices must 'switch sides'.

Formally, a *reduction rule* that is applied to a constraint $\mathcal{C}$ associated with a node α in the search tree associates α with a new constraint $\mathcal{C}'$ or rejects $\mathcal{C}$; the reduction rule is *correct* either if $\mathcal{C}$ has a fulfilling partition if and only if $\mathcal{C}'$ does, or if the rule rejects $\mathcal{C}$, then no valid monopolar partition of G fulfills $\mathcal{C}$. A *branching rule* applied to a constraint $\mathcal{C}$ associated with a node α in the search tree produces more than one child node of α, each associated with a constraint; the branching rule is *correct* if $\mathcal{C}$ has a fulfilling partition if and only if at least one of the child nodes of α is associated with a constraint $\mathcal{C}'$ that has a fulfilling partition.

The algorithm first performs the reduction rules exhaustively, in order, and then performs the branching rules, in order. That is, Reduction Rule i may only be applied if Reduction Rule i' for all $i' < i$ cannot be applied. In particular, after Reduction Rule i is applied, we start over and apply Reduction Rule 1, and so on. The same principle applies to the branching rules; moreover, branching rules are only applied if no reduction rule can be applied.

Let $\mathcal{C} = (A^{\mathcal{C}}_*, A^{\mathcal{C}}_P, B^{\mathcal{C}}_*, B^{\mathcal{C}}_P)$ be a constraint. We now describe the reduction rules. Throughout, recall from Fact 2.1 that cluster graphs contain no P_3 as an induced subgraph. The first reduction rule identifies obvious cases in which a constraint cannot be fulfilled.

Reduction Rule 5.1. If $G[A^{\mathcal{C}}_P]$ is not a cluster graph with at most k clusters, or if $G[B^{\mathcal{C}}_P]$ is not an edgeless graph, then reject the current constraint.

Proof of correctness. If $G[A^{\mathcal{C}}_P]$ is not a cluster graph with at most k clusters, then there is no valid monopolar partition (A, B) satisfying $A^{\mathcal{C}}_P \subseteq A$. Similarly, there is no valid monopolar partition (A, B) satisfying $B^{\mathcal{C}}_P \subseteq B$ if $G[B^{\mathcal{C}}_P]$ is not an edgeless graph. □

The second reduction rule finds vertices that must be moved from $B^{\mathcal{C}}_*$ to $A^{\mathcal{C}}_P$.

Reduction Rule 5.2. If there is a vertex $u \in B^{\mathcal{C}}_*$ that has a neighbor in $B^{\mathcal{C}}_P$, then set $A^{\mathcal{C}}_P \leftarrow A^{\mathcal{C}}_P \cup \{u\}$ and $B^{\mathcal{C}}_* \leftarrow B^{\mathcal{C}}_* \setminus \{u\}$; that is, replace $\mathcal{C}$ with the constraint $(A^{\mathcal{C}}_*, A^{\mathcal{C}}_P \cup \{u\}, B^{\mathcal{C}}_* \setminus \{u\}, B^{\mathcal{C}}_P)$.

Proof of correctness. For every partition (A, B) fulfilling $\mathcal{C}$, $G[B]$ is an edgeless graph and $B^{\mathcal{C}}_P \subseteq B$. Hence, $u \in A$. □

The third reduction rule finds vertices that must be moved from $A^{\mathcal{C}}_*$ to $B^{\mathcal{C}}_P$.

Reduction Rule 5.3. If there is a vertex $u \in A^{\mathcal{C}}_*$ and two vertices $w, x \in A^{\mathcal{C}}_P$ such that $G[\{u, w, x\}]$ is a P_3, set $A^{\mathcal{C}}_* \leftarrow A^{\mathcal{C}}_* \setminus \{u\}$ and $B^{\mathcal{C}}_P \leftarrow B^{\mathcal{C}}_P \cup \{u\}$.

Proof of correctness. For every partition (A, B) fulfilling $\mathcal{C}$, the graph $G[A]$ is a cluster graph and $A^{\mathcal{C}}_P \subseteq A$. Hence, $u \in B$. □

The first branching rule identifies pairs of vertices from $A^{\mathcal{C}}_*$ such that at least one of them must be moved to $B^{\mathcal{C}}_P$ because they form a P_3 with a vertex in $A^{\mathcal{C}}_P$.

Branching Rule 5.1. If there are two vertices $u, w \in A_*^{\mathcal{C}}$ and a vertex $x \in A_P^{\mathcal{C}}$ such that $G[\{u, w, x\}]$ is a P_3, then branch into two branches: one associated with the constraint $(A_*^{\mathcal{C}} \setminus \{u\}, A_P^{\mathcal{C}}, B_*^{\mathcal{C}}, B_P^{\mathcal{C}} \cup \{u\})$ and one associated with the constraint $(A_*^{\mathcal{C}} \setminus \{w\}, A_P^{\mathcal{C}}, B_*^{\mathcal{C}}, B_P^{\mathcal{C}} \cup \{w\})$.

Proof of correctness. For every partition (A, B) fulfilling $\mathcal{C}$, $G[A]$ is a cluster graph and $A_P^{\mathcal{C}} \subseteq A$. Hence, $u \in B$ or $w \in B$. □

It is important to observe that if none of the previous rules applies, then $(A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}, B_*^{\mathcal{C}} \cup B_P^{\mathcal{C}})$ is a monopolar partition (we prove this rigorously in Lemma 5.1). However, $G[A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}]$ may consist of too many clusters for this to be a *valid* monopolar partition. To check whether it is possible to reduce the number of clusters in $G[A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}]$, we apply a second branching rule that deals with singleton clusters in $G[A']$.

Branching Rule 5.2. If there is a vertex $u \in A_*^{\mathcal{C}}$ such that $\{u\}$ is a cluster in $G[A']$, then branch into two branches: the first is associated with the constraint $(A_*^{\mathcal{C}} \setminus \{u\}, A_P^{\mathcal{C}} \cup \{u\}, B_*^{\mathcal{C}}, B_P^{\mathcal{C}})$, and the second is associated with the constraint $(A_*^{\mathcal{C}} \setminus \{u\}, A_P^{\mathcal{C}}, B_*^{\mathcal{C}}, B_P^{\mathcal{C}} \cup \{u\})$.

Proof of correctness. For every partition (A, B) fulfilling $\mathcal{C}$, we have either $u \in A$ or $u \in B$. □

If no more rules apply to a constraint $\mathcal{C}$, then we can determine whether $\mathcal{C}$ can be fulfilled:

Lemma 5.1. Let $\mathcal{C} = (A_*^{\mathcal{C}}, A_P^{\mathcal{C}}, B_*^{\mathcal{C}}, B_P^{\mathcal{C}})$ be a constraint to which Reduction Rules 5.1, 5.2, and 5.3, and Branching Rules 5.1 and 5.2 do not apply. Then $(A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}, B_*^{\mathcal{C}} \cup B_P^{\mathcal{C}})$ is a monopolar partition. Moreover, there is a valid monopolar partition (A, B) fulfilling $\mathcal{C}$ if and only if $(A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}, B_*^{\mathcal{C}} \cup B_P^{\mathcal{C}})$ is valid.

Proof. First, we show that $(A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}, B_*^{\mathcal{C}} \cup B_P^{\mathcal{C}})$ is a monopolar partition. There are no induced P_3's in $G[A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}]$, because Reduction Rules 5.1 and 5.3 and Branching Rule 5.1 do not apply, and because there are no induced P_3's in G containing three vertices from $A_*^{\mathcal{C}} \subseteq A'$. Similarly, there are no edges in $G[B_*^{\mathcal{C}} \cup B_P^{\mathcal{C}}]$, because Reduction Rules 5.1 and 5.2 do not apply, and because there are no edges in $G[B']$ and $B_*^{\mathcal{C}} \subseteq B'$.

To show the second statement in the lemma, observe that, if $(A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}, B_*^{\mathcal{C}} \cup B_P^{\mathcal{C}})$ is valid, then $\mathcal{C}$ is fulfilled by $(A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}, B_*^{\mathcal{C}} \cup B_P^{\mathcal{C}})$. It remains to show that, if $(A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}, B_*^{\mathcal{C}} \cup B_P^{\mathcal{C}})$ is not valid, then each monopolar partition (A, B) of G fulfilling $\mathcal{C}$ is not valid. For the sake of contradiction, assume that this is not the case and let (A, B) be a valid monopolar partition fulfilling $\mathcal{C}$. Since $(A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}, B_*^{\mathcal{C}} \cup B_P^{\mathcal{C}})$ is a monopolar partition of G that is not valid, there are more than k clusters in $G[A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}]$. Thus, there is a cluster Q in $G[A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}]$ such that $Q \subseteq B$. Note that $|Q| = 1$, because $G[B]$ has no edges and $Q \subseteq B$. Because (A, B) fulfills $\mathcal{C}$, $Q \cap A_P^{\mathcal{C}} = \emptyset$ and thus $Q \subseteq A_*^{\mathcal{C}}$. Hence, Q is a subset of a cluster Q' of $G[A']$, as $Q \subseteq A_*^{\mathcal{C}} \subseteq A'$. However, $|Q'| \geq 2$, because Branching Rule 5.2 does not apply even though $Q \subseteq A_*^{\mathcal{C}}$. Hence, any rule that moved the

vertices of $Q' \setminus Q$ was not Branching Rule 5.2. Then the description of the other rules implies that $Q' \setminus Q \subseteq B_P^{\mathcal{C}}$. Note that $B_P^{\mathcal{C}} \subseteq B$, because (A, B) fulfills $\mathcal{C}$. Hence, $Q' \subseteq B$ and thus $G[B]$ contains an edge. Therefore, (A, B) is not a monopolar partition, a contradiction to our choice of (A, B). □

The following lemmas will be used to upper bound the depth of the search tree, and the number of applications of each rule along each root-leaf path in this tree. Herein a *leaf* of the search tree is a node associated either with a constraint that Reduction Rule 5.1 rejects, or with a constraint to which no rule applies.

Lemma 5.2. Along any root-leaf path in the search tree of the algorithm, Reduction Rule 5.2 is applied at most $k + 1$ times.

Proof. Let $\mathcal{C} = (A_*^{\mathcal{C}}, A_P^{\mathcal{C}}, B_*^{\mathcal{C}}, B_P^{\mathcal{C}})$ be a constraint obtained from an initial constraint via $k+1$ applications of Reduction Rule 5.2 and an arbitrary number of applications of Reduction Rules 5.1 and 5.3, and Branching Rules 5.1 and 5.2. Each application of Reduction Rule 5.2 adds a vertex of B' to $A_P^{\mathcal{C}}$. Since $G[B']$ is edgeless, any monopolar partition (A, B) with $A_P^{\mathcal{C}} \subseteq A$ has at least $k + 1$ clusters in $G[A]$ and, therefore, is not valid. Reduction Rule 5.1 will then be applied before any further application of Reduction Rule 5.2, and the constraint $\mathcal{C}$ will be rejected. □

Lemma 5.3. Along any root-leaf path in the search tree of the algorithm, Reduction Rule 5.3 and Branching Rules 5.1 and 5.2 are applied at most $k + 1$ times in total.

Proof. Let $\mathcal{C} = (A_*^{\mathcal{C}}, A_P^{\mathcal{C}}, B_*^{\mathcal{C}}, B_P^{\mathcal{C}})$ be a constraint obtained from an initial constraint via $k + 1$ applications of Reduction Rule 5.3 and Branching Rules 5.1 and 5.2, and an arbitrary number of applications of the other rules. Let k_s denote the number of singleton clusters in $G[A']$. Observe that each application of Reduction Rule 5.3 or Branching Rules 5.1 and 5.2 makes a vertex of $A_*^{\mathcal{C}} \subseteq A'$ permanent by placing it in $A_P^{\mathcal{C}}$ or $B_P^{\mathcal{C}}$. By the description of all rules, a vertex will never be made permanent twice. Hence, out of the $k+1$ applications of Reduction Rule 5.3 and Branching Rules 5.1 and 5.2, at most k_s make the vertex from a singleton cluster of $G[A']$ permanent. Observe that Branching Rule 5.2 cannot make a vertex from a nonsingleton cluster in $G[A']$ permanent. Thus, Reduction Rule 5.3 and Branching Rule 5.1 make at least $k - k_s + 1$ vertices in the $k - k_s$ nonsingleton clusters of $G[A']$ permanent. Since $k - k_s + 1 \geq 1$, this also implies that a nonsingleton cluster exists. By the pigeonhole principle, out of the $k - k_s + 1$ vertices that are made permanent by Reduction Rule 5.3 and Branching Rule 5.1, two are from the same nonsingleton cluster in $G[A']$. Since both Reduction Rule 5.3 and Branching Rule 5.1 only move vertices from $A_*^{\mathcal{C}}$ to $B_P^{\mathcal{C}}$, it follows that $B_P^{\mathcal{C}}$ contains two adjacent vertices. Then the constraint $\mathcal{C}$ will be rejected by Reduction Rule 5.1, which is applied before any further rule is applied. □

Theorem 5.1. Inductive Monopolar Recognition can be solved in $\mathcal{O}(2^k \cdot k^3 \cdot \Delta)$ time, where Δ is the maximum degree of G.

Proof. We call a leaf of the search tree associated with a constraint to which no rule applies an *exhausted leaf*. By Lemma 5.1 and the correctness of the rules, G has a valid monopolar partition if and only if for at least one exhausted leaf node, the partition $(A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P, B^{\mathcal{C}}_* \cup B^{\mathcal{C}}_P)$, induced by the constraint $\mathcal{C}$ associated with that node, is a valid monopolar partition. Hence, if the search tree has an exhausted leaf for which the partition $(A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P, B^{\mathcal{C}}_* \cup B^{\mathcal{C}}_P)$, induced by the constraint $\mathcal{C}$ associated with that node, is a valid monopolar partition, the algorithm answers 'yes'; otherwise, it answers 'no'. Therefore, the described search-tree algorithm correctly decides an instance of Inductive Monopolar Recognition.

To upper bound the running time, let $\mathcal{T}$ denote the search tree of the algorithm. By Lemma 5.3, Branching Rules 5.1 and 5.2 are applied at most $k+1$ times in total along any root-leaf path in $\mathcal{T}$. It follows that the depth of $\mathcal{T}$ is at most $k+2$. As each of the branching rules is a two-way branch, $\mathcal{T}$ is a binary tree, and thus the number of leaves in $\mathcal{T}$ is $\mathcal{O}(2^k)$.

The running time along any root-leaf path in $\mathcal{T}$ is dominated by the overall time taken along the path to test the applicability of the reduction and branching rules, and to apply them. By Lemma 5.2 and Lemma 5.3, along any root-leaf path in $\mathcal{T}$ the total number of applications of Reduction Rules 5.2 and 5.3 and Branching Rules 5.1 and 5.2 is $\mathcal{O}(k)$. Reduction Rule 5.1 is applied once before the application of each of the aforementioned rules. It follows that the total number of applications of all rules along any root-leaf path in $\mathcal{T}$ is $\mathcal{O}(k)$. Moreover, $\mathcal{T}$ has $\mathcal{O}(2^k)$ leaves as argued before. Therefore, we test for the applicability of the rules and apply them, or use the check of Lemma 5.1, at most $\mathcal{O}(2^k \cdot k)$ times.

We now upper bound the time to test the applicability of the rules and to apply them by $\mathcal{O}(k^2 \cdot \Delta)$. Let $\mathcal{C} = (A^{\mathcal{C}}_*, A^{\mathcal{C}}_P, B^{\mathcal{C}}_*, B^{\mathcal{C}}_P)$ be a constraint associated with a node in $\mathcal{T}$. Observe that each cluster in $G[A^{\mathcal{C}}_*]$ has size $\mathcal{O}(\Delta)$. Since $G[A^{\mathcal{C}}_*]$ has at most k clusters, this implies that $|A^{\mathcal{C}}_*| \leq k \cdot \Delta$. Thus, in $\mathcal{O}(k \cdot \Delta)$ time, we can compute a list of all clusters in $G[A^{\mathcal{C}}_*]$ and the size of each cluster. The same holds for $G[A']$. Observe that we can always check in $\mathcal{O}(1)$ time, for a given vertex v, whether v is contained in A', $A^{\mathcal{C}}_*$, $A^{\mathcal{C}}_P$, $B^{\mathcal{C}}_*$, or $B^{\mathcal{C}}_P$ and, in case v is contained in A' or $A^{\mathcal{C}}_*$, we can find the index and the size of the cluster that contains v. Moreover, by Lemma 5.2 and 5.3, we can assume that $|A^{\mathcal{C}}_P| = \mathcal{O}(k)$, and by Lemma 5.3, we can assume that $|B^{\mathcal{C}}_P| = \mathcal{O}(k)$.

To test the applicability of Reduction Rules 5.1 and 5.2, we check whether $G[A^{\mathcal{C}}_P]$ is a cluster graph with at most k clusters, whether $G[B^{\mathcal{C}}_P]$ is edgeless, and whether there is an edge with one endpoint in $B^{\mathcal{C}}_P$ and the other endpoint in $B^{\mathcal{C}}_*$. This can be done in $\mathcal{O}(k \cdot \Delta)$ time since $|A^{\mathcal{C}}_P| = \mathcal{O}(k)$ and $|B^{\mathcal{C}}_P| = \mathcal{O}(k)$.

To test the applicability of Reduction Rule 5.3, we consider each pair v, w of vertices in $A^{\mathcal{C}}_P$. If v and w are adjacent, then in $\mathcal{O}(\Delta)$ time we can check whether there is a vertex $u \in A^{\mathcal{C}}_*$ such that u is adjacent to exactly one of v and w. If v and w are not adjacent, then in $\mathcal{O}(\Delta)$ time we can check whether they have a common neighbor in $A^{\mathcal{C}}_*$. If neither condition applies to any pair v, w, then Reduction Rule 5.3 does not apply. Overall, this test takes $\mathcal{O}(k^2 \cdot \Delta)$ time.

To test the applicability of Branching Rule 5.1, we can check for each vertex

v of the at most k vertices of $A_P^{\mathcal{C}}$ in $\mathcal{O}(\Delta)$ time whether v has neighbors in two different clusters of $A_*^{\mathcal{C}}$, or whether there are two vertices u, w in the same cluster of $A_*^{\mathcal{C}}$ such that v is adjacent to u but not adjacent to w. If one of the two cases applies to some vertex $v \in A_P^{\mathcal{C}}$, then Branching Rule 5.1 applies to v. Otherwise, there is no P_3 containing exactly one vertex from $A_P^{\mathcal{C}}$ and exactly two vertices from $A_*^{\mathcal{C}}$, and Branching Rule 5.1 does not apply. Hence, the applicability of Branching Rule 5.1 can be tested in $\mathcal{O}(k \cdot \Delta)$ time.

To test the applicability of Branching Rule 5.2, we can check in $\mathcal{O}(k)$ time, whether $G[A_*^{\mathcal{C}}]$ contains a singleton cluster that is also a singleton cluster of $G[A']$.

All rules can trivially be applied in $\mathcal{O}(1)$ time if they were found to be applicable. Hence, the running time to test and apply any of the rules is $\mathcal{O}(k^2 \cdot \Delta)$.

Finally, if none of the rules applies, then we can check in $\mathcal{O}(k \cdot \Delta)$ time whether the number of clusters in $G[A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}]$ is at most k. Hence, the algorithm runs in $\mathcal{O}(2^k \cdot k^3 \cdot \Delta)$ time in total. □

Given the above theorem, Corollary 3.2 immediately implies Theorem 1.1, which we restate below:

Theorem 1.1. *In $\mathcal{O}(2^k \cdot k^3 \cdot (n+m))$ time, we can decide whether G admits a monopolar partition (A, B) such that $G[A]$ is a cluster graph with at most k clusters.*

6 Generalizations of the algorithm for Monopolar Recognition

In this section, we present two general FPTalgorithms for a range of cases of (Π_A, Π_B)-Recognition in which Π_A and Π_B are characterized by a finite set of forbidden induced subgraphs. We achieve this by adapting the algorithm of Section 5, meaning that the obtained algorithms rely on the inductive recognition framework. Therefore, the main step is to solve Inductive (Π_A, Π_B)-Recognition, where we are given a graph G with a distinguished vertex v, and a (Π_A, Π_B)-partition (A', B') of $G - v$, and we are asked to determine whether G has a (Π_A, Π_B)-partition (A, B). To solve Inductive (Π_A, Π_B)-Recognition, we consider again constraints $\mathcal{C}$ of the type $(A_*^{\mathcal{C}}, A_P^{\mathcal{C}}, B_*^{\mathcal{C}}, B_P^{\mathcal{C}})$. That is, $\mathcal{C}$ is a four-partition of the vertex set such that $A_*^{\mathcal{C}} \subseteq A'$ and $B_*^{\mathcal{C}} \subseteq B'$, where $A_P^{\mathcal{C}}$ and $B_P^{\mathcal{C}}$ represent the permanent vertices, which may not be moved between A and B anymore. We start with the two initial constraints $(A', \{v\}, B', \emptyset)$ and $(A', \emptyset, B', \{v\})$, and recursively search for solutions fulfilling one of the two constraints, building a search tree whose nodes correspond to constraints. As in Section 5, a (Π_A, Π_B)-partition (A, B) of G fulfills a constraint $\mathcal{C}$ if $A_P^{\mathcal{C}} \subseteq A$ and $B_P^{\mathcal{C}} \subseteq B$.

The first step towards designing both algorithms is to generalize several reduction and branching rules of the algorithm for Monopolar Recognition. The generalization of Reduction Rule 5.1 is as follows.

Reduction Rule 6.1. If $G[A_P^{\mathcal{C}}]$ does not fulfill Π_A or if $G[B_P^{\mathcal{C}}]$ does not fulfill Π_B, then reject the current constraint.

Proof of correctness. If $G[A_P^{\mathcal{C}}] \notin \Pi_A$, then since Π_A is hereditary, there is no A such that $A_P^{\mathcal{C}} \subseteq A$ and $G[A] \in \Pi_A$. Similarly, if $G[B_P^{\mathcal{C}}] \notin \Pi_B$, then there is no B such that $B_P^{\mathcal{C}} \subseteq B$ and $G[B] \in \Pi_B$. For any (Π_A, Π_B)-partition (A, B) fulfilling $\mathcal{C}$, we have, however, $A_P^{\mathcal{C}} \subseteq A$ and $B_P^{\mathcal{C}} \subseteq B$. Thus, no such (Π_A, Π_B)-partition exists. □

The other reduction and branching rules are aimed at destroying forbidden induced subgraphs in the candidate vertex set $A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}$ for A and in the candidate vertex set $B_*^{\mathcal{C}} \cup B_P^{\mathcal{C}}$ for B by moving vertices between $A_*^{\mathcal{C}}$ and $B_*^{\mathcal{C}}$. To destroy the forbidden induced subgraphs of Π_A in $A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}$, we used Reduction Rule 5.3 and Branching Rule 5.1 in Section 5. A generalized variant of these rules is as follows.

Branching Rule 6.1. If there is a vertex set $\tilde{A} \subseteq A_*^{\mathcal{C}} \cup A_P^{\mathcal{C}}$ such that $\tilde{A}$ is minimal under the property that $G[\tilde{A}]$ is a forbidden induced subgraph of Π_A, then for each $u \in \tilde{A} \setminus A_P^{\mathcal{C}}$ branch into a branch associated with the constraint $(A_*^{\mathcal{C}} \setminus \{u\}, A_P^{\mathcal{C}}, B_*^{\mathcal{C}}, B_P^{\mathcal{C}} \cup \{u\})$.

Proof of correctness. Suppose that (A, B) is a (Π_A, Π_B)-partition of G fulfilling $\mathcal{C}$. Since $G[\tilde{A}]$ does not fulfill Π_A, there is a vertex $u \in \tilde{A}$ such that $u \in B$. Moreover, since (A, B) fulfills $\mathcal{C}$, we have $A_P^{\mathcal{C}} \subseteq A$, and hence $A_P^{\mathcal{C}} \cap B = \emptyset$. Therefore, $u \notin A_P^{\mathcal{C}}$. Consequently, in the branch of Branching Rule 6.1 which is associated with the constraint $(A_*^{\mathcal{C}} \setminus \{u\}, A_P^{\mathcal{C}}, B_*^{\mathcal{C}}, B_P^{\mathcal{C}} \cup \{u\})$, this constraint is fulfilled by (A, B) since $u \in B$. □

In the case of Monopolar Recognition, for subgraphs that do not fulfill Π_B, it was sufficient to use Reduction Rule 5.2. This can be generalized to the following branching rule; the correctness proof is analogous to that of Branching Rule 6.1 and omitted.

Branching Rule 6.2. If there is a vertex set $\tilde{B} \subseteq B_*^{\mathcal{C}} \cup B_P^{\mathcal{C}}$ such that $\tilde{B}$ is minimal under the property that $G[\tilde{B}]$ is a forbidden induced subgraph of Π_B, then for each $u \in \tilde{B} \setminus B_P^{\mathcal{C}}$ branch into a branch associated with the constraint $(A_*^{\mathcal{C}}, A_P^{\mathcal{C}} \cup \{u\}, B_*^{\mathcal{C}} \setminus \{u\}, B_P^{\mathcal{C}})$.

In the following two subsections we use the above reduction and branching rules and some specialized rule to give the promised algorithms.

6.1 Cluster Graphs and Graphs Excluding Large Cliques

For monopolar graphs, Π_A is the family of cluster graphs with at most k cliques, and Π_B is the property of being edgeless. We now consider the more general case where Π_B excludes some clique and has a characterization by minimal forbidden induced subgraphs, each of order at most r (and Π_A remains the family of all cluster graphs with at most k cliques).

The algorithm uses Reduction Rule 6.1, Reduction Rule 5.3, Branching Rule 5.1, and Branching Rule 6.2. (It does not use Branching Rule 6.1 which is used the next subsection.) Note that applying Reduction Rule 5.3 and Branching Rule 5.1 is almost equivalent to applying Branching Rule 6.1. The difference is that we do not branch on forbidden induced subgraphs that are an edgeless graph on $k+1$ vertices. This improves the efficiency of the resulting algorithm. The following fact is not hard to prove.

Fact 6.1. Let $\mathcal{C} = (A^{\mathcal{C}}_*, A^{\mathcal{C}}_P, B^{\mathcal{C}}_*, B^{\mathcal{C}}_P)$ be such that none of Reduction Rule 6.1, Reduction Rule 5.3, Branching Rule 5.1, and Branching Rule 6.2 apply. Then, $G[A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P]$ is a cluster graph and $G[B^{\mathcal{C}}_* \cup B^{\mathcal{C}}_P]$ satisfies Π_B.

That $G[A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P]$ is a cluster graph can be seen by inapplicability of Reduction Rule 6.1, Reduction Rule 5.3, Branching Rule 5.1 and the fact that none of the rules puts any vertex into $A^{\mathcal{C}}_*$ which is initially a subset of A' of the (Π_A, Π_B)-partition (A', B'). That $G[B^{\mathcal{C}}_* \cup B^{\mathcal{C}}_P]$ satisfies Π_B follows from the inapplicability of Reduction Rule 6.1 and Branching Rule 6.2.

To obtain a (Π_A, Π_B)-partition $(A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P, B^{\mathcal{C}}_* \cup B^{\mathcal{C}}_P)$, it remains to ensure that $G[A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P]$ contains at most k clusters. If $G[A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P]$ has more than k clusters, then some vertex has to be moved from $A^{\mathcal{C}}_*$ to $B^{\mathcal{C}}_P$. This is done in the following branching rule.

Branching Rule 6.3. If $G[A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P]$ is a cluster graph with more than k clusters, then let $u \in A^{\mathcal{C}}_*$ be a vertex contained in a cluster of $G[A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P]$, such that this cluster contains no vertices from $A^{\mathcal{C}}_P$ (such a cluster must exist by Reduction Rule 6.1). Branch into two branches: one associated with the constraint $(A^{\mathcal{C}}_* \setminus \{u\}, A^{\mathcal{C}}_P \cup \{u\}, B^{\mathcal{C}}_*, B^{\mathcal{C}}_P)$ and one associated with the constraint $(A^{\mathcal{C}}_* \setminus \{u\}, A^{\mathcal{C}}_P, B^{\mathcal{C}}_*, B^{\mathcal{C}}_P \cup \{u\})$.

The rule is trivially correct since u is either contained in A or in B for any (Π_A, Π_B)-partition (A, B) fulfilling $\mathcal{C}$. Now if none of the rules applies, then we have found a solution.

Fact 6.2. Let $\mathcal{C} = (A^{\mathcal{C}}_*, A^{\mathcal{C}}_P, B^{\mathcal{C}}_*, B^{\mathcal{C}}_P)$ be such that Reduction Rule 6.1, Reduction Rule 5.3, Branching Rule 5.1, Branching Rule 6.2 do not apply to $\mathcal{C}$. Then $(A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P, B^{\mathcal{C}}_* \cup B^{\mathcal{C}}_P)$ is a (Π_A, Π_B)-partition.

To bound the running time, in particular, the number of applications of the branching and reduction rules, we make use of the fact that the properties Π_A and Π_B are mutually d-exclusive for some integer d.

Lemma 6.1. Let Π_A be the set of cluster graphs with at most k cliques, and let Π_B be any hereditary graph property that excludes the complete graph on s vertices as an induced subgraph. Then, Π_A and Π_B are mutually $((s-1)\cdot k+1)$-exclusive.

Proof. Let G be a graph of order at least $(s-1)\cdot k+1$ that fulfills Π_A, that is, G is a cluster graph with at most k clusters. By the pigeonhole principle, one of these clusters has at least s vertices. Therefore, G contains a clique on s vertices. Thus, G does not fulfill Π_B. □

We can now conclude with the complete algorithm and its running-time analysis.

Theorem 6.1. Let Π_A be the set of all cluster graphs with at most k cliques, and let Π_B be any hereditary graph property such that

- Π_B can be characterized by forbidden induced subgraphs, each of order at most r; and
- for some $s \leq r$, the complete graph on s vertices is a forbidden induced subgraph of Π_B.

Then (Π_A, Π_B)-Recognition can be solved in $2^{s \cdot k} \cdot (r-1)^{(s-1) \cdot k} \cdot n^{\mathcal{O}(1)}$ time.

Proof. The algorithm creates the two initial constraints $(A^{\mathcal{C}}_*, \{v\}, B^{\mathcal{C}}_*, \emptyset)$ and $(A^{\mathcal{C}}_*, \emptyset, B^{\mathcal{C}}_*, \{v\})$. Then, for each initial constraint, it applies Reduction Rule 6.1, Reduction Rule 5.3, Branching Rule 5.1, Branching Rule 6.2, and Branching Rule 6.3 exhaustively. If none of these rules applies to the current constraint, then the algorithm outputs $(A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P, B^{\mathcal{C}}_* \cup B^{\mathcal{C}}_P)$, which, by Fact 6.2, is a (Π_A, Π_B)-partition of G. Thus, to prove the correctness of the algorithm, it remains to show that if there is a (Π_A, Π_B)-partition (A, B) for G, then the algorithm outputs such a partition.

Suppose that (A, B) is a (Π_A, Π_B)-partition. Then (A, B) fulfills one of the initial constraints. If a constraint $\mathcal{C}$ is fulfilled by (A, B), then Reduction Rule 6.1 does not apply to this constraint, and, by the correctness of the rules, any application of Reduction Rule 5.3, Branching Rule 5.1, Branching Rule 6.2, or Branching Rule 6.3 yields at least one constraint that is fulfilled by (A, B). Hence, the initial constraint $\mathcal{C}$ fulfilling (A, B) has at least one descendant that is fulfilled by (A, B), and to which none of the reduction and branching rules applies. For this constraint, the algorithm outputs a (Π_A, Π_B)-partition.

It remains to bound the running time of the algorithm. Since all minimal forbidden induced subgraphs of Π_B have at most r vertices, we can check in $n^{\mathcal{O}(1)}$ time, whether any of the branching and reduction rules applies. To obtain the running time bound, it is thus sufficient to bound the number of created constraints in the search tree.

To this end, we bound the number of applications of the branching and reduction rules along any path from an initial constraint to a leaf constraint. Reduction Rule 6.1 is applied at most once. To bound the number of applications of the other rules, we use that, by Lemma 6.1, Π_A and Π_B are mutually $((s-1) \cdot k+1)$ exclusive. This implies that Branching Rule 6.2 is applied at most $(s-1) \cdot k+1$ times: Each application of the rule adds a vertex of B' to $A^{\mathcal{C}}_P$. Since $G[B']$ fulfills Π_B, so does $G[A^{\mathcal{C}}_P \cap B']$. Thus, if $|A^{\mathcal{C}}_P \cap B'| > (s-1) \cdot k$, then $G[A^{\mathcal{C}}_P]$ does not fulfill Π_A and Reduction Rule 6.1 applies, terminating the current branch.

Reduction Rule 5.3, Branching Rule 5.1, and Branching Rule 6.3 can be applied altogether at most $k + (s-1) \cdot k + 2$ times: Each application of any of these rules either adds a vertex of A' to $B^{\mathcal{C}}_P$ or increases the number of clusters in $G[A^{\mathcal{C}}_P]$ by one. Again by Lemma 6.1, at most $(s-1) \cdot k$ vertices of A' can be

moved from $A^{\mathcal{C}}_*$ to $B^{\mathcal{C}}_P$, before Reduction Rule 6.1 applies. Similarly, if $G[A^{\mathcal{C}}_P]$ has more than k clusters, then Reduction Rule 6.1 applies.

To bound the number of leaf constraints, observe that Branching Rule 6.2 branches into at most $r-1$ new constraints, and Branching Rule 5.1 and Branching Rule 6.3 branch into 2 new constraints. Thus, the overall number of leaf constraints is $\mathcal{O}(2^{s\cdot k+2} \cdot (r-1)^{(s-1)\cdot k+1})$. By the bound on the number of constraints on any root-leaf path, we thus have that the overall search tree size is $\mathcal{O}(2^{s\cdot k} \cdot (r-1)^{(s-1)\cdot k} \cdot k^{\mathcal{O}(1)})$ which implies the overall running time bound. □

6.2 Mutually Exclusive Graph Properties with Small Forbidden Subgraphs

We now relax the demands on the hereditary properties Π_A and Π_B even further: we demand only that Π_A excludes some fixed edgeless graph, that Π_B excludes some fixed clique, and that Π_A and Π_B are each characterized by forbidden induced subgraphs of constant size. Recall that, by Proposition 4.1, such properties are mutually d-exclusive for some constant d.

Theorem 6.2. Let Π_A and Π_B be two hereditary graph properties such that

- Π_A excludes an edgeless graph of order c_A and has a characterization by forbidden induced subgraphs, each of order at most r_A; and
- Π_B excludes a complete graph of order c_B and has a characterization by forbidden induced subgraphs of order at most r_B.

Then (Π_A, Π_B)-Recognition can be solved in $(r_A-1)^{R(c_A,c_B)} \cdot (r_B-1)^{R(c_A,c_B)} \cdot n^{\mathcal{O}(1)}$ time.

Proof. We show that Inductive (Π_A, Π_B)-Recognition can be solved in this running time. In conjunction with Theorem 3.1, this implies the above Theorem 6.2.

The algorithm performs a search from the two initial constraints of Fact 5.2. For each constraint encountered during the search, we check if Reduction Rule 6.1 applies. If this is not the case, we check if Branching Rule 6.1 or Branching Rule 6.2 applies. If neither applies, then $(A^{\mathcal{C}}_* \cup A^{\mathcal{C}}_P, B^{\mathcal{C}}_* \cup B^{\mathcal{C}}_P)$ is a (Π_A, Π_B)-partition. Otherwise, apply the respective rule and continue the search with the constraints created by the rule. By the correctness of Branching Rule 6.1, Branching Rule 6.2, and Fact 5.2, this algorithm finds a (Π_A, Π_B)-partition of G if it exists.

It remains to analyze the running time of the algorithm. First, observe that throughout the algorithm, we have $A^{\mathcal{C}}_* \subseteq A'$ and $B^{\mathcal{C}}_* \subseteq B'$. Thus, every vertex set $\tilde{A}$ to which Branching Rule 6.1 applies contains at least one vertex from $A^{\mathcal{C}}_P$ and therefore creates at most $r_A - 1$ new recursive branches. Now, observe that each application of Branching Rule 6.1 increases the number of vertices in $B^{\mathcal{C}}_P$ by one. Moreover, all vertices in $B^{\mathcal{C}}_P$ are from A'. Thus, if $|B^{\mathcal{C}}_P| \geq R(c_A, c_B)$, then by Proposition 4.1, this implies that $G[B^{\mathcal{C}}_P] \notin \Pi_B$. Thus, Branching Rule 6.1 is applied at most $R(c_A, c_B)$ times before Reduction

Rule 6.1 applies. Similarly, Branching Rule 6.2 is applied at most $R(c_A, c_B)$ times before Reduction Rule 6.1 applies and each application of Branching Rule 6.2 creates at most $r_B - 1$ constraints. Overall, the number of created constraints is thus $\mathcal{O}((r_A - 1)^{R(c_A,c_B)} \cdot (r_B - 1)^{R(c_A,c_B)})$. For each constraint, we must check if any of the reduction or branching rules applies, which can be done in $n^{\mathcal{O}(1)}$ time by using the polynomial-time algorithms for checking membership of a graph in Π_A and Π_B. □

For the special case of recognizing monopolar graphs with at most k cliques, Theorem 6.2 applies: The graphs fulfilling Π_A have at most k clusters, thus the edgeless graph of order $k+1$ is forbidden, implying $c_A = k+1$. The forbidden subgraphs for Π_A are exactly the P_3 and the edgeless graph on $k+1$ vertices, implying $r_A = k+1$. For Π_B, the only forbidden subgraph is the clique on two vertices, implying $c_B = r_B = 2$. Altogether, this gives a running time of $k^{R(k+1,2)} \cdot n^{O(1)} = k^k \cdot n^{O(1)}$. Hence, our tailored algorithm in Section 5 is substantially more efficient than the generic algorithm.

Another application of Theorem 6.2 is to (Π_A, Π_B)-Recognition for Π_A being the triangle-free graphs and Π_B being the complete graphs. The running time becomes $\mathcal{O}(2^{R(3,2)} \cdot nm)$ in this case, yielding an $O(nm)$-time algorithm with very small constant hidden in the $\mathcal{O}$-Notation.

7 An FPT Algorithm for 2-Subcoloring

In this section, we give an FPT algorithm for 2-Subcoloring parameterized by the smaller number of clusters in the two parts. Although the general approach is similar to the approach used for Monopolar Recognition, in that it relies on inductive recognition and the notion of constraints, the algorithm is substantially more complex. In particular, the notion of constraints and the reduction and branching rules are more involved.

Throughout, given a graph G and a nonnegative integer k, we call a 2-subcoloring (A, B) of G *valid* if $G[A]$ has at most k cliques. In the inductive recognition framework, we need a parameterized inductive recognizer for the following problem:

> Inductive 2-Subcoloring
> **Input:** A graph $G = (V, E)$, a vertex $v \in V$, and a valid 2-subcoloring (A', B') of $G' = G - v$.
> **Question:** Does G have a valid 2-subcoloring (A, B)?

Fix an instance of Inductive 2-Subcoloring with a graph $G = (V, E)$, a vertex $v \in V$, and a valid 2-subcoloring (A', B') of $G' = G - v$. We again apply a search-tree algorithm that starts with initial partitions $(A_*^{\mathcal{C}}, B_*^{\mathcal{C}})$ of V, derived from (A', B'), that are not necessarily 2-subcolorings of G. Then, we try to "repair" those partitions by moving vertices between $A_*^{\mathcal{C}}$ and $B_*^{\mathcal{C}}$ to form a valid 2-subcoloring (A, B) of G. As before, each node in the search tree is associated with one constraint.

Definition 7.1. A *constraint* $\mathcal{C} = (A_1^{\mathcal{C}}, \ldots, A_k^{\mathcal{C}}, B_1^{\mathcal{C}}, \ldots, B_n^{\mathcal{C}}, A_P^{\mathcal{C}}, B_P^{\mathcal{C}})$ consists of a partition $(A_1^{\mathcal{C}}, \ldots, A_k^{\mathcal{C}}, B_1^{\mathcal{C}}, \ldots, B_n^{\mathcal{C}})$ of V and two vertex sets $A_P^{\mathcal{C}} \subseteq A_*^{\mathcal{C}}$ and $B_P^{\mathcal{C}} \subseteq B_*^{\mathcal{C}}$, where $A_*^{\mathcal{C}} = \bigcup_{i=1}^{k} A_i^{\mathcal{C}}$ and $B_*^{\mathcal{C}} = \bigcup_{i=1}^{n} B_i^{\mathcal{C}}$, such that for any $i \neq j$:

- u and w are not adjacent for any $u \in A_i^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$ and $w \in A_j^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$, and
- u and w are not adjacent for any $u \in B_i^{\mathcal{C}} \setminus B_P^{\mathcal{C}}$ and $w \in B_j^{\mathcal{C}} \setminus B_P^{\mathcal{C}}$.

We explicitly allow (some of) the sets of the partition $(A_1^{\mathcal{C}}, \ldots, A_k^{\mathcal{C}}, B_1^{\mathcal{C}}, \ldots, B_n^{\mathcal{C}})$ of V to be empty. The vertices in $A_P^{\mathcal{C}}$ and $B_P^{\mathcal{C}}$ are called *permanent* vertices of the constraint.

The permanent vertices in $A_P^{\mathcal{C}}$ and $B_P^{\mathcal{C}}$ in the definition will correspond precisely to those vertices that have switched sides during the algorithm. We refer to the sets $A_1^{\mathcal{C}}, \ldots, A_k^{\mathcal{C}}$ and $B_1^{\mathcal{C}}, \ldots, B_n^{\mathcal{C}}$ as *groups*; during the algorithm, $G[A_*^{\mathcal{C}}]$ and $G[B_*^{\mathcal{C}}]$ are not necessarily cluster graphs and, thus, we avoid using the term clusters.

We now define the notion of a valid 2-subcoloring fulfilling a constraint. Intuitively speaking, a constraint $\mathcal{C}$ is fulfilled by a bipartition (A, B) if (A, B) respects the assignment of the permanent vertices stipulated by $\mathcal{C}$, and if all vertices that do not switch sides stay in the bipartition (A, B) in the same groups they belong to in $\mathcal{C}$. This notion is formalized as follows.

Definition 7.2. A constraint $\mathcal{C} = (A_1^{\mathcal{C}}, \ldots, A_k^{\mathcal{C}}, B_1^{\mathcal{C}}, \ldots, B_n^{\mathcal{C}}, A_P^{\mathcal{C}}, B_P^{\mathcal{C}})$ is *fulfilled* by a bipartition (A, B) of V if $G[A]$ is a cluster graph with k clusters $A_1, \ldots, A_k$ and $G[B]$ is a cluster graph with n clusters $B_1, \ldots, B_n$ (some of the clusters may be empty) such that:

1. for each $i \in [k]$, $A_i \cap A_*^{\mathcal{C}} \subseteq A_i^{\mathcal{C}}$;
2. for each $i \in [n]$, $B_i \cap B_*^{\mathcal{C}} \subseteq B_i^{\mathcal{C}}$;
3. $A_P^{\mathcal{C}} \subseteq A$; and
4. $B_P^{\mathcal{C}} \subseteq B$.

We now need a set of initial constraints to jumpstart the search-tree algorithm.

Lemma 7.1. Let $A'_1, \ldots, A'_k$ denote the clusters of $G'[A']$ and let $B'_1, \ldots, B'_n$ denote the clusters of $G'[B']$. Herein, if there are less than k clusters in $G[A']$ or less than n clusters in $G[B']$, we add an appropriate number of empty sets. By relabeling, we may assume that only $B'_1, \ldots, B'_i$ contain neighbors of v, and $B'_{i+1} = \emptyset$. Each valid 2-subcoloring (A, B) of G fulfills either:

- $(A'_1, \ldots, A'_j \cup \{v\}, \ldots, A'_k, B'_1, \ldots, B'_n, \{v\}, \emptyset)$ for some $j \in [k]$, or
- $(A'_1, \ldots, A'_k, B'_1, \ldots, B'_j \cup \{v\}, \ldots, B'_n, \emptyset, \{v\})$ for some $j \in [i+1]$.

Proof. Since (A', B') is a valid 2-subcoloring of $G' = G - v$ and $v \in A_P^{\mathcal{C}} \cup B_P^{\mathcal{C}}$ for each constraint $\mathcal{C}$, the constructed tuples are indeed constraints. Let $A_1, \ldots, A_k$ be the clusters of $G[A]$ and $B_1, \ldots, B_n$ be those of $G[B]$.

First, assume that $v \in A$. If, for some $j \in [k]$, there is a vertex $u \in A'_j \cap N(v) \cap A$, then (A, B) fulfills the constraint $(A'_1, \ldots, A'_j \cup \{v\}, \ldots, B'_n, \{v\}, \emptyset)$. This can be seen as follows. Since $G[A]$ is a cluster graph, each A_i contains vertices of at most one cluster of $(A'_1, \ldots, A'_k)$. Similarly, each B_i contains vertices of at most one cluster of $(B'_1, \ldots, B'_n)$. Hence, the clusters of $G[A]$ and $G[B]$ can be labeled accordingly to satisfy Conditions 1 and 2 of Definition 7.2. Moreover, by assumption, $\{v\} \subseteq A$ and thus Conditions 3 and 4 of Definition 7.2 are fulfilled.

Similarly, if, for all $j \in [k]$, vertex v has no neighbors in $A'_j \cap A$, then there is a $j \in [k]$ such that $A'_j \cap A = \emptyset$. Then, (A, B) fulfills the constraint $(A'_1, \ldots, A'_j \cup \{v\}, \ldots, B'_n, \{v\}, \emptyset)$, by the same arguments as above.

Symmetric arguments apply for the case $v \in B$. □

Now that we have identified the initial constraints, we turn to the search-tree algorithm and its reduction and branching rules. A crucial ingredient to the rules and the analysis of the running time is the following lemma. A consequence of the lemma is that if the number of initial constraints is too large, then most of them should be rejected immediately.

Lemma 7.2. Let $\mathcal{C} = (A^{\mathcal{C}}_1, \ldots, A^{\mathcal{C}}_k, B^{\mathcal{C}}_1, \ldots, B^{\mathcal{C}}_n, A^{\mathcal{C}}_P, B^{\mathcal{C}}_P)$ be a constraint and let (A, B) be any valid 2-subcoloring of G fulfilling $\mathcal{C}$. If $u \in V$ has neighbors in more than $k+1$ groups among $B^{\mathcal{C}}_1, \ldots, B^{\mathcal{C}}_n$, then $u \in A$.

Proof. For the sake of contradiction, suppose that there is a valid 2-subcoloring (A, B) of G fulfilling $\mathcal{C}$ such that $u \in B$. Let $A_1, \ldots, A_k$ and $B_1, \ldots, B_n$ denote the (possibly empty) clusters of $G[A]$ and $G[B]$ respectively. Let $b_1, \ldots, b_{k+2}$ be neighbors of u in distinct groups among $B^{\mathcal{C}}_1, \ldots, B^{\mathcal{C}}_n$; these vertices exist by assumption. Without loss of generality, $b_i \in B^{\mathcal{C}}_i$ for each $i = 1, \ldots, k+2$.

Since (A, B) fulfills $\mathcal{C}$, we have $B_i \cap B^{\mathcal{C}}_* \subseteq B^{\mathcal{C}}_i$ for each $i = 1, \ldots, k+2$. Therefore, the vertices among $b_1, \ldots, b_{k+2}$ that are in B are in distinct clusters of $G[B]$. Since $u \in B$, at least $k+1$ vertices among $b_1, \ldots, b_{k+2}$ are in A, say $b_1, \ldots, b_{k+1} \in A$. Since (A, B) fulfills $\mathcal{C}$, we have $b_1, \ldots, b_{k+1} \notin B^{\mathcal{C}}_P$, and, thus, $b_i \in B^{\mathcal{C}}_i \setminus B^{\mathcal{C}}_P$ for each $i = 1, \ldots, k+1$. Then the definition of a constraint implies that $b_1, \ldots, b_{k+1}$ are pairwise nonadjacent. Hence, these vertices are in distinct clusters of $G[A]$. Therefore, $G[A]$ has at least $k+1$ clusters, a contradiction. □

Lemma 7.2 implies that if v has neighbors in more than $k+1$ clusters of A', then we should immediately reject the initial constraints generated by Lemma 7.1 that place v in $B^{\mathcal{C}}_P$. Hence, we obtain the following corollary of Lemmas 7.1 and 7.2.

Corollary 7.1. Lemma 7.1 generates at most $2k+2$ constraints that are not immediately rejected.

Proof. Lemma 7.2 implies that, if v has neighbors in more than $k+1$ clusters of B', then we should immediately reject the initial constraints generated by Lemma 7.1 that place v in $B^{\mathcal{C}}_P$. □

As before, each nonroot node of the search tree is associated with a constraint. The root of the search tree is a dummy node with children associated with the constraints generated by Lemma 7.1 that are not immediately rejected due to Lemma 7.2. We now give two reduction rules, which are applied exhaustively to each search-tree node, in the order they are presented. Let $\mathcal{C} = (A^{\mathcal{C}}_1, \ldots, A^{\mathcal{C}}_k, B^{\mathcal{C}}_1, \ldots, B^{\mathcal{C}}_n, A^{\mathcal{C}}_P, B^{\mathcal{C}}_P)$ be a constraint. The first reduction rule identifies some obvious cases in which the constraint cannot be fulfilled.

Reduction Rule 7.1. If $G[A_P^{\mathcal{C}}]$ or $G[B_P^{\mathcal{C}}]$ is not a cluster graph, or if there are $i \neq j$ such that there is an edge between $A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}}$ and $A_j^{\mathcal{C}} \cap A_P^{\mathcal{C}}$ or an edge between $B_i^{\mathcal{C}} \cap B_P^{\mathcal{C}}$ and $B_j^{\mathcal{C}} \cap B_P^{\mathcal{C}}$, then reject $\mathcal{C}$.

Proof of correctness. If $G[A_P^{\mathcal{C}}]$ is not a cluster graph, then there is no valid 2-subcoloring (A, B) such that $A_P^{\mathcal{C}} \subseteq A$. Similarly, there is no valid 2-subcoloring (A, B) such that $B_P^{\mathcal{C}} \subseteq B$ if $G[B_P^{\mathcal{C}}]$ is not a cluster graph.

If there is an edge between two vertices $u \in A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}}$ and $w \in A_j^{\mathcal{C}} \cap A_P^{\mathcal{C}}$ where $i \neq j$, then every valid 2-subcoloring (A, B) needs to have u and w in the same cluster of $G[A]$, which contradicts Condition 1. Similarly, Condition 2 is violated if there is an edge between two vertices $u \in B_i^{\mathcal{C}} \cap B_P^{\mathcal{C}}$ and $w \in B_j^{\mathcal{C}} \cap B_P^{\mathcal{C}}$ where $i \neq j$. □

The second reduction rule is the natural consequence of Lemma 7.2.

Reduction Rule 7.2. If there is a vertex $u \in A_i^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$ that has neighbors in more than $k+1$ groups of $B_*^{\mathcal{C}}$, then set $A_P^{\mathcal{C}} \leftarrow A_P^{\mathcal{C}} \cup \{u\}$.

The algorithm contains a single branching rule. This rule, called switch(u), uses branching to fix a vertex u in one of the clusters in one of the parts of the 2-subcoloring. The vertices to which switch() must be applied are identified by *switching rules*. We say that a switching rule that calls for applying switch(u) is *correct* if for all valid 2-subcolorings (A, B) of G fulfilling $\mathcal{C}$, we have $u \in A_*^{\mathcal{C}} \cap B$ or $u \in B_*^{\mathcal{C}} \cap A$. We first describe the switching rules, and then describe switch(u).

Recall from Fact 2.1 that cluster graphs do not contain induced P_3's.

The next rule identifies vertices that are not adjacent to some permanent vertices of their group.

Switching Rule 7.1. If there is a vertex u such that $u \in A_i^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$ and u is not adjacent to some vertex in $A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}}$, or $u \in B_i^{\mathcal{C}} \setminus B_P^{\mathcal{C}}$ and u is not adjacent to some vertex in $B_i^{\mathcal{C}} \cap B_P^{\mathcal{C}}$, then call switch($u$).

Proof of correctness. Let (A, B) be a valid 2-subcoloring fulfilling $\mathcal{C}$ and let $(A_1, \ldots, A_k)$ be the partition of A induced by the clusters of $G[A]$. We show the correctness of the case $u \in A_i^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$ (the other case is symmetric). Suppose that $u \in A$. By Condition 1, $w \in A_i$ for any $w \in A_i^{\mathcal{C}} \cap A$. Hence, $u \in A_i$. Moreover, since $A_P^{\mathcal{C}} \subseteq A$ by Condition 3, $A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}} \subseteq A_i^{\mathcal{C}} \cap A$ and, thus, $w \in A_i$ for any $w \in A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}}$. However, $G[\{u\} \cup (A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}})]$ is not a clique by assumption, contradicting that A_i is a clique. □

The final switching rule finds vertices that have permanent neighbors in another group.

Switching Rule 7.2. If there is a vertex u such that $u \in A_i^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$ and u has a neighbor in $A_P^{\mathcal{C}} \setminus A_i^{\mathcal{C}}$, or $u \in B_i^{\mathcal{C}} \setminus B_P^{\mathcal{C}}$ and u has a neighbor in $B_P^{\mathcal{C}} \setminus B_i^{\mathcal{C}}$, then call switch($u$).

Proof of correctness. Let (A, B) be a valid 2-subcoloring fulfilling $\mathcal{C}$. We show the correctness of the case $u \in A_i^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$ (the other case is symmetric). If $u \in A$, then u must be in the same cluster of $G[A]$ as its neighbor in $A_P^{\mathcal{C}} \setminus A_i^{\mathcal{C}}$, because $A_P^{\mathcal{C}} \subseteq A$ by Condition 3. However, this contradicts Condition 1. □

Now, we describe switch(u), which is a combination of a reduction rule and a branching rule. There are two main scenarios that we distinguish. If u has permanent neighbors in the other part, then there is only one choice for assigning u to a group. Otherwise, we branch into all (up to symmetry when a group is empty) possibilities to place u into a group. It is important to note that the switching rules never apply switch(u) to a permanent vertex.

Branching Rule 7.1 (switch(u))**.**

- If $u \in A_i^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$ and u has a permanent neighbor in some $B_j^{\mathcal{C}}$, then set $A_i^{\mathcal{C}} \leftarrow A_i^{\mathcal{C}} \setminus \{u\}$, $B_j^{\mathcal{C}} \leftarrow B_j^{\mathcal{C}} \cup \{u\}$, $B_P^{\mathcal{C}} \leftarrow B_P^{\mathcal{C}} \cup \{u\}$.
- If $u \in A_i^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$ and u has only nonpermanent neighbors in $B_*^{\mathcal{C}}$, then, for each $B_j^{\mathcal{C}}$ such that $N(u) \cap B_j^{\mathcal{C}} \neq \emptyset$ and $B_j^{\mathcal{C}} \cap B_P^{\mathcal{C}} = \emptyset$, and for one $B_j^{\mathcal{C}}$ such that $B_j^{\mathcal{C}} = \emptyset$ (chosen arbitrarily), branch into a branch associated with the constraint $(A_1^{\mathcal{C}}, \ldots, A_i^{\mathcal{C}} \setminus \{u\}, \ldots, A_k^{\mathcal{C}}, B_1^{\mathcal{C}}, \ldots, B_j^{\mathcal{C}} \cup \{u\}, \ldots, B_n^{\mathcal{C}}, A_P^{\mathcal{C}}, B_P^{\mathcal{C}} \cup \{u\})$.
- If $u \in B_i^{\mathcal{C}} \setminus B_P^{\mathcal{C}}$ and u has a permanent neighbor in some $A_j^{\mathcal{C}}$, then set $B_i^{\mathcal{C}} \leftarrow B_i^{\mathcal{C}} \setminus \{u\}$, $A_j^{\mathcal{C}} \leftarrow A_j^{\mathcal{C}} \cup \{u\}$, $A_P^{\mathcal{C}} \leftarrow A_P^{\mathcal{C}} \cup \{u\}$.
- If $u \in B_i^{\mathcal{C}} \setminus B_P^{\mathcal{C}}$ and u has only nonpermanent neighbors in $A_*^{\mathcal{C}}$, then for each $A_j^{\mathcal{C}}$ with $A_j^{\mathcal{C}} \cap A_P^{\mathcal{C}} = \emptyset$, branch into a branch associated with the constraint $(A_1^{\mathcal{C}}, \ldots, A_j^{\mathcal{C}} \cup \{u\}, \ldots, A_k^{\mathcal{C}}, B_1^{\mathcal{C}}, \ldots, B_i^{\mathcal{C}} \setminus \{u\}, \ldots, B_n^{\mathcal{C}}, A_P^{\mathcal{C}} \cup \{u\}, B_P^{\mathcal{C}})$; if no such $A_j^{\mathcal{C}}$ exists, reject $\mathcal{C}$.

Proof of correctness. All tuples produced by Branching Rule 7.1 are indeed constraints, since u is made permanent in every case. Let (A, B) be a valid 2-subcoloring fulfilling $\mathcal{C}$. We prove that (A, B) fulfills one of the constraints created by Branching Rule 7.1. We consider two cases.

Case 1: $u \in A_i^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$. Since the switching rule calling switch(u) is correct, we have $u \in B$ and, thus, $B_P^{\mathcal{C}} \cup \{u\} \subseteq B$. Thus, for all constraints produced by switch(u), the 2-subcoloring (A, B) satisfies Conditions 3 and 4 of Definition 7.2. Furthermore, Condition 1 is satisfied in every generated constraint because removing u from $A_*^{\mathcal{C}}$ weakens the requirement placed on A. It remains to show that (A, B) satisfies Condition 2 for one of the generated constraints. If u has a neighbor $w \in B_P^{\mathcal{C}}$, then u is in the same cluster as w in $G[B]$ since (A, B) fulfills $\mathcal{C}$ and by Condition 2 of fulfilling constraints. Hence, for the (single) constraint produced by switch(u) in this situation, the 2-subcoloring (A, B) satisfies Condition 2. Now consider the case that u has no neighbor in $B_P^{\mathcal{C}}$. Let $B_1, \ldots, B_n$ be the clusters in $G[B]$, indexed according to the groups in $\mathcal{C}$ and padded with empty sets if the number of clusters is smaller than n. Let B_j be the cluster that contains u. As u does not have neighbors in $B_P^{\mathcal{C}}$, cluster B_j

does not contain any vertex of $B_P^{\mathcal{C}}$. There are two subcases. First, u has a neighbor $w \in B_*^{\mathcal{C}}$. As (A, B) fulfills $\mathcal{C}$ and by Condition 2, we have $w \in B_j^{\mathcal{C}}$. Thus, (A, B) satisfies Condition 2 for the constraint generated by switch(u) in which u is added to $B_j^{\mathcal{C}}$. Second, u does not have a neighbor in $B_*^{\mathcal{C}}$. Hence, $B_j \cap B_*^{\mathcal{C}} = \emptyset$ and hence, without loss of generality, by relabeling we have $B_j^{\mathcal{C}} = \emptyset$. Therefore, for the constraint in which u is added to $B_j^{\mathcal{C}} = \emptyset$, the 2-subcoloring (A, B) satisfies Condition 2.

Case 2: $u \in B_i^{\mathcal{C}} \setminus B_P^{\mathcal{C}}$. In this case the argument is simpler. Since the switching rule invoking switch(u) is correct, we have $u \in A$, and thus $A_P^{\mathcal{C}} \cup \{u\} \subseteq A$. This means that, for all constraints produced by switch(u), the 2-subcoloring (A, B) fulfills Conditions 3 and 4. The case that u has a neighbor in $w \in A_P^{\mathcal{C}}$ follows by an argument symmetric to the one of Case 1. If u has no neighbor $w \in A_P^{\mathcal{C}}$, then switch(u) considers all possibilities of placing u in one of the clusters of A. Again, u cannot be in a cluster in $G[A]$ that contains a vertex from $A_P^{\mathcal{C}}$, meaning that the groups containing permanent vertices may be ignored in the branching. Hence, for one of the produced constraints, (A, B) fulfills Conditions 1 and 2. However, special consideration is needed if each of $A_1^{\mathcal{C}}, \ldots, A_k^{\mathcal{C}}$ has a permanent vertex. Let (A, B) be a valid 2-subcoloring of G that fulfills $\mathcal{C}$. Then Condition 1 and 3 imply that $G[A]$ has k clusters that each contains a permanent vertex of $A_P^{\mathcal{C}}$. However, u is not adjacent to any permanent vertices. Hence, $u \notin A$, which contradicts the correctness of the switching rule that called switch(u). Therefore, $\mathcal{C}$ cannot be fulfilled, and the rule correctly rejects $\mathcal{C}$. □

If none of the previous rules applies, then the constraint directly gives a solution:

Lemma 7.3. Let $\mathcal{C} = (A_1^{\mathcal{C}}, \ldots, B_n^{\mathcal{C}}, A_P^{\mathcal{C}}, B_P^{\mathcal{C}})$ be a constraint such that none of the rules applies. Then $(A_*^{\mathcal{C}}, B_*^{\mathcal{C}})$ is a valid 2-subcoloring.

Proof. We need to show that $G[A_*^{\mathcal{C}}]$ and $G[B_*^{\mathcal{C}}]$ are cluster graphs and that $G[A_*^{\mathcal{C}}]$ has at most k clusters. First, we claim that $G[A_i^{\mathcal{C}}]$ is a clique for every $i = 1, \ldots, k$. Every vertex in $A_i^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$ is adjacent to every vertex in $A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}}$; otherwise, Switching Rule 7.1 applies. Any two vertices in $A_i^{\mathcal{C}} \setminus A_P^{\mathcal{C}}$ are also adjacent, because they are in the same cluster of A'. It remains to show that $G[A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}}]$ is a clique. By the description of switch(u), if a vertex x is placed into $A_i^{\mathcal{C}}$ and $A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}} \neq \emptyset$, then x is adjacent to a vertex of $A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}}$. Hence, $G[A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}}]$ is connected. Since Reduction Rule 7.1 does not apply, $G[A_i^{\mathcal{C}} \cap A_P^{\mathcal{C}}]$ does not contain an induced P_3 and, thus, it is a clique. Hence, $G[A_i^{\mathcal{C}}]$ is a clique, as claimed.

Second, we claim that there are no edges between $A_i^{\mathcal{C}}$ and $A_j^{\mathcal{C}}$, where $i \neq j$. Suppose for the sake of a contradiction that e is such an edge. Since Reduction Rule 7.1 does not apply, e is incident with at least one nonpermanent vertex. Since Switching Rule 7.2 does not apply, e is in fact incident with two nonpermanent vertices. Then e cannot exist by the definition of a constraint. The claim follows.

The combination of the above claims shows that $G[A^{\mathcal{C}}_*]$ is a cluster graph with the clusters $A^{\mathcal{C}}_i$ (some of which may be empty) and, thus, has at most k clusters. Similar arguments show that $G[B^{\mathcal{C}}_*]$ is a cluster graph: in the above argument, we used only Reduction Rule 7.1 and Switching Rules 7.1 and 7.2, which apply to vertices in $A^{\mathcal{C}}_*$ and $B^{\mathcal{C}}_*$ symmetrically. □

Using the above rules and lemmas, we can now show the following.

Theorem 7.1. Inductive 2-Subcoloring can be solved in $\mathcal{O}(k^{2k+1} \cdot (n+m))$ time.

Proof. Given the valid 2-subcoloring (A', B') of G', we use Lemma 7.1 to generate a set of initial constraints, and reject those which cannot be fulfilled due to Lemma 7.2. By Corollary 7.1, at most $2k + 2$ initial constraints remain, which are associated with the children of the (dummy) root node. For each node of the search tree, we first exhaustively apply the reduction rules on the associated constraint. Afterwards, if there exists a vertex u to which a switching rule applies, then we apply switch(u). If switch(u) does not branch but instead reduces to a new constraint, then we apply the reduction rules exhaustively again, etc.

A *leaf* of the search tree is a node associated either with a constraint that is rejected, or with a constraint to which no rule applies. The latter is called an *exhausted leaf*. If the search tree has an exhausted leaf, then the algorithm answers 'yes'; otherwise, it answers 'no'. By the correctness of the reduction, branching, and switching rules, and by Lemma 7.3, graph G has a valid 2-subcoloring if and only if the search tree has at least one exhausted leaf node. Therefore, the described search-tree algorithm correctly decides an instance of Inductive 2-Subcoloring.

We now bound the running time of the algorithm. Observe that each described reduction rule and the branching rule switch() either rejects the constraint or makes a vertex permanent. Hence, along each root-leaf path, $\mathcal{O}(n)$ rules are applied. Each rule can trivially be tested for applicability and applied in polynomial time. Hence, it remains to bound the number of leaves of the search tree.

As mentioned, at the root of the search tree, we create at most $\mathcal{O}(n)$ constraints, out of which at most $2k + 2$ constraints do not correspond to leaf nodes by Lemma 7.1, Corollary 7.1 and Reduction Rule 7.2. The only branches are created by a call to switch(u) for a vertex u that has only nonpermanent neighbors in the other part of the bipartition $(A^{\mathcal{C}}_*, B^{\mathcal{C}}_*)$. Observe that if such a vertex $u \in B^{\mathcal{C}}_* \setminus B^{\mathcal{C}}_P$, then in each constraint $\mathcal{C}'$ constructed by switch(u) the number of groups in $A^{\mathcal{C}'}_*$ that have at least one permanent vertex increases by one compared to $\mathcal{C}$. Since each constraint has k groups in $A^{\mathcal{C}}_*$, this branch can be applied at most k times along each root-leaf path in the search tree.

Similarly, if $u \in A^{\mathcal{C}}_* \setminus A^{\mathcal{C}}_P$, then in each constraint $\mathcal{C}'$ constructed by switch(u) the number of groups in $B^{\mathcal{C}}_*$ that have at least one permanent vertex increases by one compared to $\mathcal{C}$. We claim that, if $B^{\mathcal{C}}_*$ has k groups with a permanent vertex, then u has a neighbor in $B^{\mathcal{C}}_P$. First, each permanent vertex in $B^{\mathcal{C}}_*$ is part of A' by the description of the rules. Moreover, the permanent vertices

of the k groups in $B^{\mathcal{C}}_*$ with a permanent vertex stem from k different clusters in $G[A']$, because switch() places a vertex of $A^{\mathcal{C}}_* \setminus A^{\mathcal{C}}_P$ that has neighbors in $B^{\mathcal{C}}_P$ in the same group as its neighbors in $B^{\mathcal{C}}_P$. This implies that one of the clusters in $G[A']$ that the permanent vertices stem from contains u. Hence, u is adjacent to a vertex in $B^{\mathcal{C}}_P$, as claimed. The claim implies that if $B^{\mathcal{C}}_*$ has k groups with a permanent vertex, then switch(u) applied to a vertex $u \in A^{\mathcal{C}}_* \setminus A^{\mathcal{C}}_P$ does not branch. Hence, also the branch of switch(u) in which $u \in A^{\mathcal{C}}_* \setminus A^{\mathcal{C}}_P$ is performed at most k times along each root-leaf path in the search tree.

In summary, the branchings of switch(u) in which $u \in B^{\mathcal{C}}_* \setminus B^{\mathcal{C}}_P$ branch into at most k cases, and the branchings in which $u \in A^{\mathcal{C}}_* \setminus A^{\mathcal{C}}_P$ branch into at most $k+2$ cases, since Reduction Rule 7.2 does not apply. Observe that k of the initial constraints have already one group in $A^{\mathcal{C}}_*$ with a permanent vertex, and the other $k+1$ initial constraints have one group in B with a permanent vertex. Thus, if the initial constraint $\mathcal{C}$ places v in $A^{\mathcal{C}}_P$, then the overall number of constraints from $\mathcal{C}$ by branching is at most $k^{k-1} \cdot (k+2)^k$. If the initial constraint $\mathcal{C}$ places v in $B^{\mathcal{C}}_P$, then the overall number of constraints created from $\mathcal{C}$ by branching is at most $k^k \cdot (k+2)^{k-1}$. Altogether, the number of constraints created by branching is thus

$$(2k+1) \cdot k^k \cdot (k+2)^k = (2k+1) \cdot k^k \cdot k^k \cdot [(1+1/(k/2))^{k/2}]^2 = \mathcal{O}(k^{2k+1})$$

after noting that $[(1+1/(k/2))^{k/2}]^2 = \mathcal{O}(1)$. This provides the claimed bound on the number of leaves of the search tree. We now give the detailed proof of the running time.

Our goal is to achieve (almost) linear running time per search-tree node. To this end, we pursue the following strategy for applying the reduction, switching, and branching rules. Note that each rule except Reduction Rule 7.2 can become applicable only due to making a vertex permanent, that is, by placing a vertex into $A^{\mathcal{C}}_P$ or $B^{\mathcal{C}}_P$. Hence, it suffices to check whether any rule applies whenever we make a vertex permanent; the applications of Reduction Rule 7.2 which are not caused by making a vertex permanent receive special treatment below. We now argue that, whenever we make a vertex u permanent, we can determine in $\mathcal{O}(k \cdot \deg(u))$ time, whether any rule applies (and a vertex to which it applies), after an initial, one-time expense of $\mathcal{O}(m+n)$ time. After this, we prove that this is enough to show the running time bound of $\mathcal{O}(k^{2k+1} \cdot (m+n))$.

First, observe that we can initialize and maintain in $\mathcal{O}(m+n)$ time throughout the algorithm a data structure that allows us to determine in $\mathcal{O}(1)$ time for an arbitrary vertex to which group it belongs and whether it is permanent, and to determine in $\mathcal{O}(1)$ time for an arbitrary group how many permanent vertices it contains. We additionally maintain a data structure that allows us to determine in $\mathcal{O}(1)$ time for an arbitrary permanent vertex to which cluster it belongs in $G[A^{\mathcal{C}}_P]$ or $G[B^{\mathcal{C}}_P]$, and to determine in $\mathcal{O}(1)$ time for an arbitrary cluster in $G[A^{\mathcal{C}}_P]$ or $G[B^{\mathcal{C}}_P]$ how many vertices it contains. This can also be done in $\mathcal{O}(m+n)$ time overall, since each vertex becomes permanent only once.

To check whether Reduction Rule 7.1 becomes applicable when we make a vertex u permanent, it suffices to check whether all permanent neighbors of u are in the same cluster and whether these neighbors include all vertices in this

cluster. This can clearly be done in $\mathcal{O}(\deg(u))$ time using the data structures mentioned above.

For Switching Rule 7.1, we first iterate over the $\deg(u)$ neighbors of u, labeling each with a "timestamp", an integer that is initially 0 and increases whenever we make a vertex permanent. Then, we iterate over a list of vertices in the group of u (note that, within overall linear time per search tree node, we can maintain these lists for all groups). For each vertex v in this list, we check in $\mathcal{O}(1)$ time whether it is labeled with the current timestamp. If not, then u and v are not adjacent and Switching Rule 7.1 applies. After iterating over at most $\deg(u)+1$ vertices, we encounter a vertex that is nonadjacent to u if there is one.

We can check for the applicability of Switching Rule 7.2 in $\mathcal{O}(\deg(u))$ time by examining each neighbor of u using the aforementioned data structures.

For Reduction Rule 7.2, note first that, except for the possible applications after the initial constraints have been created, Reduction Rule 7.2 can only become applicable when we move a vertex to $B^{\mathcal{C}}_*$. Whenever we move a vertex to $B^{\mathcal{C}}_*$ in any of the rules, we also make it permanent. Except for the initial applications, which we treat below, it thus suffices to check whether Reduction Rule 7.2 becomes applicable whenever we make a vertex u permanent. To do this in $\mathcal{O}(k \cdot \deg(u))$ time, we maintain throughout the algorithm for each nonpermanent vertex in $A^{\mathcal{C}}_*$ a list with at most k entries containing the indices of the groups in $B^{\mathcal{C}}_*$ in which it has neighbors. This list can be initialized in $\mathcal{O}(m+n)$ time in the beginning. Whenever we move a vertex u to $A^{\mathcal{C}}_*$ (and make it permanent), we initialize such a list for u. Whenever we move a vertex u to $B^{\mathcal{C}}_*$ (and make it permanent), we update the list for each neighbor. Both cases take $\mathcal{O}(k \cdot \deg(u))$ time. At the same time, we can check whether Reduction Rule 7.2 becomes applicable and identify the vertex to which it becomes applicable.

Concluding, whenever we make a vertex u permanent, we can determine in $\mathcal{O}(k \cdot \deg(u))$ time whether any rule becomes applicable, and, if so, find a corresponding vertex to apply it to if needed. We now show that this suffices to prove an overall running time of $\mathcal{O}(k^{2k+1} \cdot (m+n))$ time.

To see this, note first that each reduction and switching rule, including the reduction part of Branching Rule 7.1, can be carried out in $\mathcal{O}(1)$ time, once we have determined whether they are applicable, and a vertex to which they are applicable if they need one. (By carrying out a switching rule, we mean to add the vertex u to a queue of vertices to which we shall apply switch().) Furthermore, we make each vertex permanent only once. Hence, the time for checking the applicability and for applying the reduction and switching rules along each root-leaf path in the search tree is bounded by $\sum_{u \in V} \mathcal{O}(k \cdot \deg(u)) = \mathcal{O}(k \cdot (m+n))$. Herein, we include the reduction part of Branching Rule 7.1 but exclude the applications of Reduction Rule 7.2 that are not due to making a vertex permanent. This running time also subsumes the time for the branching part of Branching Rule 7.1 by attributing the time taken for constructing constraints in a search-tree node to its child nodes.

To finish the proof, it remains to treat Reduction Rule 7.2 and the pruning of the initial constraints generated by Lemma 7.2. For Reduction Rule 7.2, observe that, after an initial exhaustive application when we place v, the only ap-

plications are due to moving a vertex between $A^{\mathcal{C}}_*$ and $B^{\mathcal{C}}_*$, which we already accounted for above. Furthermore, Reduction Rule 7.2 can be exhaustively applied in $\mathcal{O}(m+n)$ time, since no vertices are moved to $B^{\mathcal{C}}_*$ in the process. Hence, the overall time needed for Reduction Rule 7.2 along a root-leaf path in the search tree is $\mathcal{O}(k\cdot(m+n))$. For the initial constraints generated by Lemma 7.2, observe that we can check in linear time whether v has more than k neighbors in groups in $B^{\mathcal{C}}_*$, and, if so, make v permanent in $A^{\mathcal{C}}_*$; then we only generate the appropriate at most k constraints. Otherwise, we can directly create the at most $2k+2$ constraints. In both cases, the initial constraints can be created in $\mathcal{O}(k\cdot(m+n))$ time. The overall time upper bound of $\mathcal{O}(k^{2k+1}\cdot(m+n))$ follows. □

Given the above theorem, we obtain our running time bound for 2-Subcoloring.

Theorem 1.3. *In $\mathcal{O}(k^{2k+1}\cdot nm)$ time, we can decide whether G admits a 2-subcoloring (A,B) such that $G[A]$ is a cluster graph with at most k clusters.*

Proof. Theorem 7.1 and Corollary 3.1 immediately imply a running time bound of $\mathcal{O}(k^{2k+1}\cdot(n^2+nm))$. Moreover, before starting inductive recognition, we may remove all vertices of degree 0 because they can be safely added to B. Afterwards, $n=\mathcal{O}(m)$, giving the claimed running time. □

8 More FPT Results

In this section, we consider two examples of parameterized (Π_A,Π_B)-Recognition problems for which the branching technique seems to outperform inductive recognition, either in terms of simplicity or efficiency. The first example we consider is that of (Π_A,Π_B)-Recognition problems, in which $|A|\leq k$, where k is the parameter, and Π_B can be characterized by a finite set of forbidden induced subgraphs. We show that this problem can be solved in time $2^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$ by a straightforward branching strategy. We note that the inductive recognition technique can be used to solve the problem within the same time upper bound, albeit with a little more work. The second example we consider is the 2-Subcoloring problem parameterized by the total number of clusters in both sides of the partition. We show that the aforementioned problem can be solved in $\mathcal{O}(4^k\cdot k^2\cdot n^2)$ time by a branching strategy, followed by reducing the resulting instances to the 2-CNF-Sat problem. We note that the algorithm in the previous section for 2-Subcoloring can be modified to solve this problem, but runs in time $\mathcal{O}(k^k nm)$.

8.1 Parameterizing (Π_A,Π_B)-Recognition by the Cardinality of A

In this subsection, we consider the parameter consisting of the total number of vertices in $G[A]$. The following proposition shows that a straightforward branching strategy yields a generic fixed-parameter algorithm for many (Π_A,Π_B)-Recognition problems:

Proposition 1.1. *Let Π_A and Π_B be two hereditary graph properties such that membership of Π_A can be decided in polynomial time and Π_B can be characterized by a finite set of forbidden induced subgraphs. Then we can decide in $2^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$ time whether V can be partitioned into sets A and B such that $G[A] \in \Pi_A$, $G[B] \in \Pi_B$, and $|A| \leq k$.*

Proof. We describe a branching algorithm. Initially, let $A = \emptyset$ and $B = V$. Now consider a branch where we are given (A, B). If $|A| > k$, then reject the current branch. Using the assumed finite set $\mathcal{H}$ of forbidden induced subgraphs that characterizes Π_B, find a forbidden induced subgraph $H = (W, F)$ in $G[B]$. This takes polynomial time, since $\mathcal{H}$ is finite. If H does not exist, then accept the current branch and answer 'yes'. Otherwise, branch into $|W|$ branches: for each $w \in W$, we construct the branch where $A' = A \cup \{w\}$ and $B' = B \setminus \{w\}$. If the algorithm never accepts a branch, then we answer 'no'. The correctness of the algorithm is straightforward: for each forbidden induced subgraph that the algorithm finds, there is a branch where the algorithm adds the vertex to A that is also in A in a solution. Every node in the search tree has finite degree, since $|W|$ is finite. Moreover, the search tree has depth at most k, since each branch adds a single vertex to A and $|A|$ cannot exceed k. Hence, the running time is $2^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$, as claimed. □

If Π_B has only an infinite characterization by forbidden subgraphs, then (Π_A, Π_B)-Recognition parameterized by $|A|$ is W[2]-hard in some cases: The problem of deleting at most k vertices in an undirected graph such that the resulting graph has no so-called wheel as induced subgraph is W[2]-hard with respect to k [21]. Thus, if Π_A is the class of graphs of order at most k and Π_B is the class of wheel-free graphs, then (Π_A, Π_B)-Recognition is W[2]-hard with respect to $|A|$.

8.2 Parameterizing 2-Subcoloring by the Total Number of Clusters

In this subsection, we prove Theorem 1.4 by presenting an FPT algorithm for 2-Subcoloring parameterized by the total number of clusters in both sides of the partition. Throughout, given a graph G and a nonnegative integer k, we call a 2-subcoloring (A, B) of G *valid* if $G[A]$ and $G[B]$ are cluster graphs that, in total, have at most k clusters.

As mentioned, the presented algorithm does not rely on the inductive recognition technique, but instead performs branching on the entire graph, followed (possibly) by reducing the resulting instances to the 2-CNF-Sat problem. To describe the branching algorithm, we define a notion of a constraint. Throughout, let k be a nonnegative integer and let $G = (V, E)$ be a graph for which we want to decide whether it has a valid 2-subcoloring.

Definition 8.1. Let $k_1, k_2 \in \mathbb{N}$. A *constraint* $\mathcal{C}$ is a partition

$$(A_1^{\mathcal{C}}, \ldots, A_{k_1}^{\mathcal{C}}, B_1^{\mathcal{C}}, \ldots, B_{k_2}^{\mathcal{C}}, R^{\mathcal{C}})$$

of V such that $G[A^{\mathcal{C}}_*]$ induces a cluster graph with the clusters $A^{\mathcal{C}}_1, \ldots, A^{\mathcal{C}}_{k_1}$, and $G[B^{\mathcal{C}}_*]$ induces a cluster graph with the clusters $B^{\mathcal{C}}_1, \ldots, B^{\mathcal{C}}_{k_2}$, where $A^{\mathcal{C}}_* = \bigcup_{i=1}^{k_1} A^{\mathcal{C}}_i$ and $B^{\mathcal{C}}_* = \bigcup_{j=1}^{k_2} B^{\mathcal{C}}_j$. We explicitly allow $k_1 = 0$ or $k_2 = 0$, meaning that there are no parts $A^{\mathcal{C}}_i$ or no parts $B^{\mathcal{C}}_i$ in these cases. A constraint is *fulfilled* by a valid 2-subcoloring (A, B) of G if $A^{\mathcal{C}}_* \subseteq A$ and $B^{\mathcal{C}}_* \subseteq B$.

The following proposition is crucial at several places in the algorithm:

Proposition 8.1. If a constraint $\mathcal{C} = (A^{\mathcal{C}}_1, \ldots, A^{\mathcal{C}}_{k_1}, B^{\mathcal{C}}_1, \ldots, B^{\mathcal{C}}_{k_2}, R^{\mathcal{C}})$ can be fulfilled by a valid 2-subcoloring (A, B) of G, then $G[A]$ contains k_1 clusters (say $A_1, \ldots, A_{k_1}$) and $G[B]$ contains k_2 clusters (say $B_1, \ldots, B_{k_2}$) such that $A^{\mathcal{C}}_i \subseteq A_i$ for each $i \in [k_1]$ and $B^{\mathcal{C}}_j \subseteq B_j$ for each $j \in [k_2]$.

Proof. Recall that $A^{\mathcal{C}}_* \subseteq A$ and $B^{\mathcal{C}}_* \subseteq B$. Moreover, observe that $u \in A^{\mathcal{C}}_i$ for $i \in [k_1]$ and $v \in A^{\mathcal{C}}_{i'}$ for $i' \in [k_1]$ are adjacent if and only if $i = i'$, because $G[A^{\mathcal{C}}_*]$ is a cluster graph with the clusters $A^{\mathcal{C}}_1, \ldots, A^{\mathcal{C}}_{k_1}$. Hence, $u \in A^{\mathcal{C}}_i$ for $i \in [k_1]$ and $v \in A^{\mathcal{C}}_{i'}$ for $i' \in [k_1]$ are in the same cluster of $G[A]$ if and only $i = i'$. A similar observation holds with respect to B. □

The algorithm is a search-tree algorithm. Each node in the search tree is associated with a constraint. The root of the search tree is associated with the constraint $\mathcal{C}_r = (R^{\mathcal{C}}_r = V)$ (that is, all vertices of G are in $R^{\mathcal{C}}_r$ and there are no clusters). At a node in the search tree, the algorithm searches for a solution that fulfills the associated constraint $\mathcal{C}$. To this end, the algorithm applies reduction and branching rules that successively tighten $\mathcal{C}$. We describe these rules next. As usual, reduction rules are performed exhaustively before branching rules, and the rules are performed in the order they are presented below. Throughout, let $\mathcal{C} = (A^{\mathcal{C}}_1, \ldots, A^{\mathcal{C}}_{k_1}, B^{\mathcal{C}}_1, \ldots, B^{\mathcal{C}}_{k_2}, R^{\mathcal{C}})$ be a constraint associated with a node of the search tree of the algorithm, and let $A^{\mathcal{C}}_* = \bigcup_{i=1}^{k_1} A^{\mathcal{C}}_i$ and $B^{\mathcal{C}}_* = \bigcup_{j=1}^{k_2} B^{\mathcal{C}}_j$.

Reduction Rule 8.1. If $k_1 + k_2 > k$, then reject $\mathcal{C}$.

Proof of correctness. By Proposition 8.1, no 2-subcoloring of G that fulfills $\mathcal{C}$ can be valid. □

For simplicity, we call $A^{\mathcal{C}}_1, \ldots, A^{\mathcal{C}}_{k_1}$ the *clusters of* $A^{\mathcal{C}}_*$ and $B^{\mathcal{C}}_1, \ldots, B^{\mathcal{C}}_{k_2}$ the *clusters of* $B^{\mathcal{C}}_*$. When we *open* a new cluster in $A^{\mathcal{C}}_*$ with a vertex $v \in R^{\mathcal{C}}$, we create the constraint $\mathcal{C}' = (A^{\mathcal{C}'}_1, \ldots, A^{\mathcal{C}'}_{k_1+1}, B^{\mathcal{C}'}_1, \ldots, B^{\mathcal{C}'}_{k_2}, R^{\mathcal{C}'})$ where $A^{\mathcal{C}'}_i = A^{\mathcal{C}}_i$ for $i \in [k_1]$, $A^{\mathcal{C}'}_{k_1+1} = \{v\}$, $B^{\mathcal{C}'}_j = B^{\mathcal{C}}_j$ for $j \in [k_2]$, and $R^{\mathcal{C}'} = R^{\mathcal{C}} \setminus \{v\}$. Opening a new cluster in $B^{\mathcal{C}}_*$ is similarly defined. When we *add* a vertex $v \in R^{\mathcal{C}}$ to $A^{\mathcal{C}}_*$, we create the constraint $\mathcal{C}'$ that differs from $\mathcal{C}$ in that $v \notin R^{\mathcal{C}'}$ but instead $v \in A^{\mathcal{C}'}_i$, where v is adjacent to all vertices of $A^{\mathcal{C}}_i$. Adding a vertex to $B^{\mathcal{C}}_*$ is similarly defined. Note that we can only open a new cluster or add a vertex to a cluster if this indeed yields a constraint; for the vertices to which these operations are applied during the algorithm, this will always be ensured.

The next reduction rule treats vertices which have to be in $A^{\mathcal{C}}_*$ or $B^{\mathcal{C}}_*$ due to connections to the corresponding clusters.

Reduction Rule 8.2. If $v \in R^{\mathcal{C}}$ is adjacent to two clusters of $A_*^{\mathcal{C}}$ (resp. $B_*^{\mathcal{C}}$) or else if v is adjacent to some but not all vertices of a cluster of $A_*^{\mathcal{C}}$ (resp. $B_*^{\mathcal{C}}$), then

- if v is adjacent to two clusters of $B_*^{\mathcal{C}}$ (resp. $A_*^{\mathcal{C}}$) or if v is adjacent to some but not all vertices of a cluster of $B_*^{\mathcal{C}}$ (resp. $A_*^{\mathcal{C}}$), then reject $\mathcal{C}$;
- if v is not adjacent to any cluster in $B_*^{\mathcal{C}}$ (resp. $A_*^{\mathcal{C}}$), then open a new cluster in $B_*^{\mathcal{C}}$ (resp. $A_*^{\mathcal{C}}$) with v;
- otherwise, add v to $B_*^{\mathcal{C}}$ (resp. to $A_*^{\mathcal{C}}$).

Proof of correctness. Suppose that $v \in R^{\mathcal{C}}$ is adjacent to two clusters of $A_*^{\mathcal{C}}$. Then in any valid 2-subcoloring (A, B) of G that fulfills $\mathcal{C}$, v is also adjacent to two clusters of A by Proposition 8.1. Hence, $v \in B$. Similarly, if v is adjacent to some but not all vertices of a cluster of $A_*^{\mathcal{C}}$, then v is also adjacent to some but not all vertices of a cluster of A by Proposition 8.1, and thus $v \in B$. By a symmetric argument, this immediately shows that the first part of the rule correctly rejects $\mathcal{C}$.

Suppose that v is not adjacent to any cluster in $B_*^{\mathcal{C}}$. Then in any valid 2-subcoloring (A, B) of G that fulfills $\mathcal{C}$ and that satisfies $v \in B$, v must be in a different cluster of $G[B]$ than the vertices of $B_*^{\mathcal{C}}$ by Proposition 8.1, because v is not adjacent to any vertex of $B_*^{\mathcal{C}}$. Hence, the second part of the rule correctly opens a new cluster with v.

If none of the previous parts of the rule applied to v, then v is adjacent to all vertices of a cluster $B_i^{\mathcal{C}}$ of $B_*^{\mathcal{C}}$. Then in any valid 2-subcoloring (A, B) of G that fulfills $\mathcal{C}$ and that satisfies $v \in B$, v must be in the same cluster as the vertices of $B_i^{\mathcal{C}}$ by Proposition 8.1. Hence, the third part of the rule correctly adds v to $B_i^{\mathcal{C}}$. □

Several types of ambiguous vertices remain. First, some of these vertices are treated by branching rules. The status of the remaining vertices will be determined via a reduction to 2-CNF-Sat. The first type of vertices have no connections to any cluster.

Branching Rule 8.1. If $v \in R^{\mathcal{C}}$ is not adjacent to any cluster of $A_*^{\mathcal{C}}$ or $B_*^{\mathcal{C}}$, then branch into two branches: in the first, open a new cluster in $A_*^{\mathcal{C}}$ with v; in the second, open a new cluster in $B_*^{\mathcal{C}}$ with v.

Proof of correctness. Let (A, B) be any valid 2-subcoloring of G that fulfills $\mathcal{C}$. Suppose that $v \in A$. Then v must be in a different cluster of $G[A]$ than the vertices of $A_*^{\mathcal{C}}$ by Proposition 8.1, because v is not adjacent to any vertex of $A_*^{\mathcal{C}}$. Hence, (A, B) fulfills the constraint generated in the first branch. Similarly, if $v \in B$, then (A, B) fulfills the constraint generated in the second branch. □

After this rule has been applied, we can obtain a useful partition of $R^{\mathcal{C}}$, classifying vertices according to the clusters in $A_*^{\mathcal{C}}$ or $B_*^{\mathcal{C}}$ to which they belong, if they are put into $A_*^{\mathcal{C}}$ or $B_*^{\mathcal{C}}$, respectively.

Proposition 8.2. If Reduction Rule 8.2 and Branching Rule 8.1 cannot be applied, then $R^{\mathcal{C}}$ can be partitioned into sets $N_{A_i^{\mathcal{C}}}$, $N_{B_j^{\mathcal{C}}}$, and $N_{A_i^{\mathcal{C}}B_j^{\mathcal{C}}}$ for $i \in [k_1]$ and $j \in [k_2]$, where:

- any vertex in $N_{A_i^{\mathcal{C}}}$ is adjacent to all vertices of cluster $A_i^{\mathcal{C}}$ and not to any other clusters of $A_*^{\mathcal{C}}$ nor any clusters of $B_*^{\mathcal{C}}$;
- any vertex in $N_{B_j^{\mathcal{C}}}$ is adjacent to all vertices of cluster $B_j^{\mathcal{C}}$ and not to any other clusters of $B_*^{\mathcal{C}}$ nor any clusters of $A_*^{\mathcal{C}}$; and
- any vertex in $N_{A_i^{\mathcal{C}}B_j^{\mathcal{C}}}$ is adjacent to all vertices of clusters $A_i^{\mathcal{C}}$ and $B_j^{\mathcal{C}}$ and not to any other clusters of $A_*^{\mathcal{C}}$ or $B_*^{\mathcal{C}}$.

Proof. By Branching Rule 8.1, any vertex $v \in R^{\mathcal{C}}$ is adjacent to at least one cluster of $A_*^{\mathcal{C}}$ or $B_*^{\mathcal{C}}$. By Reduction Rule 8.2, v cannot be adjacent to two clusters of $A_*^{\mathcal{C}}$ or two clusters of $B_*^{\mathcal{C}}$. Furthermore, again by Reduction Rule 8.2, if v is adjacent to a cluster $A_i^{\mathcal{C}}$ of $A_*^{\mathcal{C}}$ (resp. $B_i^{\mathcal{C}}$ of $B_*^{\mathcal{C}}$), then v is adjacent to all vertices of that cluster. □

This partition of $R^{\mathcal{C}}$ enables further branching rules. The following two branching rules take care of pairs of vertices that cannot belong to the same cluster in one of the parts $A_*^{\mathcal{C}}$, $B_*^{\mathcal{C}}$.

Branching Rule 8.2. If $u, v \in N_{A_i^{\mathcal{C}}}$ for some $i \in [k_1]$ (resp. $u, v \in N_{B_j^{\mathcal{C}}}$ for some $j \in [k_2]$) and $uv \notin E$, then branch into two branches: in the first, open a new cluster in $B_*^{\mathcal{C}}$ (resp. $A_*^{\mathcal{C}}$) with u; in the second, open a new cluster in $B_*^{\mathcal{C}}$ (resp. $A_*^{\mathcal{C}}$) with v.

Proof of correctness. Let (A, B) be any valid 2-subcoloring of G that fulfills $\mathcal{C}$. Suppose $u, v \in N_{A_i^{\mathcal{C}}}$ for some $i \in [k_1]$. Then $A_i^{\mathcal{C}} \subseteq A_i$ for some cluster A_i of $G[A]$ by Proposition 8.1. Since $uv \notin E$, at least one of u, v is not in A_i. In fact, since u and v are adjacent to all vertices of $A_i^{\mathcal{C}}$ by definition, at least one of u, v is not in A. By definition, u and v are not adjacent to any vertex of $B_*^{\mathcal{C}}$. Hence, at least one of u, v must be in a different cluster of $G[B]$ than the vertices of $B_*^{\mathcal{C}}$ by Proposition 8.1. Therefore, (A, B) fulfills the generated constraint in at least one of the branches. The argument in case $u, v \in N_{B_j^{\mathcal{C}}}$ for some $j \in [k_2]$ follows symmetrically. □

Branching Rule 8.3. If $u \in N_{A_i^{\mathcal{C}}}$ and $v \in N_{A_{i'}^{\mathcal{C}}}$ for some $i, i' \in [k_1]$ with $i \neq i'$ (resp. $u \in N_{B_j^{\mathcal{C}}}$ and $v \in N_{B_{j'}^{\mathcal{C}}}$ for some $j, j' \in [k_2]$ with $j \neq j'$) and $uv \in E$, then branch into two branches: in the first, open a new cluster in $B_*^{\mathcal{C}}$ (resp. $A_*^{\mathcal{C}}$) with u; in the second, open a new cluster in $B_*^{\mathcal{C}}$ (resp. $A_*^{\mathcal{C}}$) with v.

Proof of correctness. Let (A, B) be any valid 2-subcoloring of G that fulfills $\mathcal{C}$. Suppose that $u \in N_{A_i^{\mathcal{C}}}$ and $v \in N_{A_{i'}^{\mathcal{C}}}$ for some $i, i' \in [k_1]$ with $i \neq i'$. Then $A_i^{\mathcal{C}} \subseteq A_i$ for some cluster A_i of $G[A]$ and $A_{i'}^{\mathcal{C}} \subseteq A_{i'}$ for some cluster $A_{i'}$ of $G[A]$ by Proposition 8.1, where $i \neq i'$. Since $uv \in E$, $u \notin A_i$ or $v \notin A_{i'}$. In fact, since

u (resp. v) is adjacent to all vertices of $A_i^{\mathcal{C}}$ (resp. $A_{i'}^{\mathcal{C}}$) by definition, at least one of u, v is not in A. By definition, u and v are not adjacent to any vertex of $B_*^{\mathcal{C}}$. Hence, at least one of u, v must be in a different cluster of $G[B]$ than the vertices of $B_*^{\mathcal{C}}$ by Proposition 8.1. Therefore, (A, B) fulfills the generated constraint in at least one of the branches. The argument in case $u \in N_{B_j^c}$ and $v \in N_{B_{j'}^c}$ for some $j, j' \in [k_2]$ with $j \neq j'$ follows symmetrically. □

Below, let $I = \{i \mid 1 \leq i \leq k_1, N_{A_i^c} \neq \emptyset\}$ and let $J = \{j \mid 1 \leq j \leq k_2, N_{B_j^c} \neq \emptyset\}$. If none of the above rules applies, then we have the following:

Proposition 8.3. If none of the above rules applies, then $G[\bigcup_{i \in I} N_{A_i^c}]$ is a cluster graph whose clusters are $N_{A_i^c}$, where $i \in I$, and $G[\bigcup_{j \in J} N_{B_j^c}]$ is a cluster graph whose clusters are $N_{B_j^c}$, where $j \in J$.

Proof. If $N_{A_i^c}$ is not a clique, then Branching Rule 8.2 applies. If a vertex of $N_{A_i^c}$ is adjacent to a vertex of $N_{A_{i'}^c}$ for $i \in I$ and $i' \in I \setminus \{i\}$, then Branching Rule 8.3 applies. The same holds mutatis mutandis with respect to $N_{B_j^c}$. □

The following branching rule is special in the sense that it will be applied only once in a root-leaf path of the corresponding search tree.

Branching Rule 8.4. Let $I = \{i \mid 1 \leq i \leq k_1, N_{A_i^c} \neq \emptyset\}$ and $J = \{j \mid 1 \leq j \leq k_2, N_{B_j^c} \neq \emptyset\}$. For each $I' \subseteq I$ and $J' \subseteq J$ such that $|I'| + |J'| \leq k - k_1 - k_2$, branch into a branch where we:

- for each $i \in I'$, open a new cluster in $B_*^{\mathcal{C}}$ with a fresh dummy vertex s_i that is made adjacent to all vertices of $N_{A_i^c}$;
- for each $j \in J'$, open a new cluster in $A_*^{\mathcal{C}}$ with a fresh dummy vertex t_j that is made adjacent to all vertices of $N_{B_j^c}$;
- for each $i \in I \setminus I'$, add all vertices of $N_{A_i^c}$ to $A_i^{\mathcal{C}}$; and
- for each $i \in J \setminus J'$, add all vertices of $N_{B_j^c}$ to $B_j^{\mathcal{C}}$.

Proof of correctness. Recall Proposition 8.3. Let (A, B) be any valid 2-subcoloring of G that fulfills $\mathcal{C}$. Let I^* denote the set containing all integers i such that B contains at least one vertex from $N(A_i^{\mathcal{C}})$. Similarly, let J^* denote the set containing all integers j such that A contains at least one vertex from $N(B_j^{\mathcal{C}})$. Each integer in $I^* \cup J^*$ corresponds to a distinct cluster that does not intersect with any cluster of the constraint $\mathcal{C}$. Thus, since (A, B) is valid, we have $|I^*| + |J^*| \leq k - k_1 - k_2$. Consequently, there is a constraint $\mathcal{C}'$ that was created in the branch where $I' = I^*$ and $J' = J^*$. Let G' denote the graph constructed for this branch, that is, the graph G plus the dummy vertices. Then, the bipartition $(A' := A \cup \{s_i \mid i \in I'\}, B' := B \cup \{t_j \mid j \in J'\})$ of G' is a solution fulfilling the constraint $\mathcal{C}'$: For each $i \in I \setminus I'$, A' contains a cluster $A_i = A_i^{\mathcal{C}}$. Moreover, for each $i \in I'$, B' contains a cluster $B_i = N(A_i^{\mathcal{C}}) \cup \{s_i\}$. The clusters

for each $j \in J$ can be defined symmetrically. Since (A, B) is a 2-subcoloring of G and by the construction of the s_i and t_j, there are no edges between different clusters in A or in B. Hence, (A', B') is a 2-subcoloring. Moreover, the bipartition (A', B') does not contain any further clusters and by the restriction on $|I'| + |J'|$ it is thus valid. □

This completes the description of the reduction and branching rules. After exhaustive application of the rules, the problem of finding a solution fulfilling the constraint associated with a node can be solved by reduction to 2-CNF-Sat.

Definition 8.2. Say that a node in the search tree is *exhausted* if none of the reduction and branching rules apply in that node.

Lemma 8.1. There is an algorithm that takes as input the graph $G = (V, E)$ and a constraint $\mathcal{C} = (A_1^{\mathcal{C}}, \ldots, A_{k_1}^{\mathcal{C}}, B_1^{\mathcal{C}}, \ldots, B_{k_2}^{\mathcal{C}}, R^{\mathcal{C}})$ associated with an exhausted leaf node, and decides in $\mathcal{O}(n^2)$ time whether G has a valid 2-subcoloring that fulfills $\mathcal{C}$.

Proof. We construct an instance of 2-CNF-Sat (satisfiability of Boolean formulas in the conjunctive normal form in which each clause contains at most two literals) in $\mathcal{O}(n^2)$ time that can be satisfied if and only if $\mathcal{C}$ can be fulfilled. Since 2-CNF-Sat is solvable in linear time (see Papadimitriou [23] for example), the algorithm runs in $O(n^2)$ time.

Observe that it suffices to place the vertices of $R^{\mathcal{C}}$ into clusters. So let v be a vertex in $R^{\mathcal{C}}$. Since no rules apply to $\mathcal{C}$ and in particular Branching Rule 8.2 and Branching Rule 8.3 do not apply, $v \in N_{A_i^c B_j^c}$ for some $i \in [k_1]$ and $j \in [k_2]$. Suppose that (A, B) is a valid 2-subcoloring that fulfills $\mathcal{C}$. It follows from Proposition 8.1 that if $v \in A$, then $v \in A_i$, and if $v \in B$, then $v \in B_j$. Therefore, what is left is to decide is whether the vertices in the sets $N_{A_i^c B_j^c}$ can be partitioned between the A and B in such a way that v can only be placed in one of its two associated clusters, and such that the resulting bipartition is a 2-subcoloring of V. We model this as an instance Φ of 2-CNF-Sat as follows.

For each vertex $v \in R^{\mathcal{C}}$, we create a Boolean variable x_v. Assigning $x_v = 1$ corresponds to adding v to its associated cluster in A, and assigning $x_v = 0$ corresponds to adding v to its associated cluster in B. The Boolean formula Φ is constructed as follows. For every two vertices $v, v' \in R^{\mathcal{C}}$, let (say) $v \in N_{A_i^c B_j^c}$ and $v' \in N_{A_{i'}^c B_{j'}^c}$. We add the clause $(\overline{x_v} \vee \overline{x_{v'}})$ to Φ if v and v' are adjacent but $i \neq i'$ or if v and v' are nonadjacent but $i = i'$; we add the clause $(x_v \vee x_{v'})$ to Φ if v and v' are adjacent but $j \neq j'$ or if v and v' are nonadjacent but $j = j'$. In both cases, the added clauses enforce that v and v' are assigned to different cluster-groups in any satisfying assignment of Φ. This completes the construction of Φ. It is not difficult to verify that $\mathcal{C}$ can be fulfilled if and only if Φ is satisfiable. Indeed, the correctness proof follows along similar lines as the correctness of Branching Rules 8.2 and 8.3.

To carry out the construction of the 2-CNF formula in $\mathcal{O}(n^2)$ time, proceed as follows. First, construct the adjacency matrix of G in $\mathcal{O}(n^2)$ time. Then, for

each vertex $v \in R^{\mathcal{C}}$ determine i and j such that $v \in N_{A_i^{\mathcal{C}} B_j^{\mathcal{C}}}$. That is, compute an array that maps each vertex v to the tuple (i, j). This can be done in $\mathcal{O}(m)$ time by iterating over all neighbors of the clusters in $A_*^{\mathcal{C}}$ and $B_*^{\mathcal{C}}$, and setting the entries of the arrays corresponding to each neighbor accordingly. Finally, iterate over all pairs of vertices in $R^{\mathcal{C}}$, determine whether they are adjacent in $\mathcal{O}(1)$ time, determine whether their i and j-values differ in $\mathcal{O}(1)$ time using the computed array, and add the clause in $\mathcal{O}(1)$ time accordingly. Hence, the algorithm runs in $\mathcal{O}(n^2)$ time, as required. □

We now combine the reduction and branching rules with Lemma 8.1 and prove the overall running-time bound.

Theorem 1.4. *In $\mathcal{O}(4^k \cdot k^2 \cdot n^2)$ time, we can decide whether G admits a 2-subcoloring (A, B) such that $G[A]$ and $G[B]$ are cluster graphs with at most k clusters in total.*

Proof. As outlined in the beginning of this section, the algorithm works as follows. We start with the constraint $\mathcal{C}_r = (R_r^{\mathcal{C}} = V)$. Then we apply the reduction and branching rules exhaustively, in the order they were presented. A leaf of the search tree is then a node of the search tree with associated constraint $\mathcal{C}$ such that either $\mathcal{C}$ is rejected or none of the rules applies to $\mathcal{C}$. The latter is an *exhausted leaf*. Lemma 8.1 shows that there is an $\mathcal{O}(n^2)$ time algorithm that decides if there is a valid 2-subcoloring of G that fulfills the constraint associated with an exhausted leaf. If this algorithm accepts the constraint associated with at least one exhausted leaf, then we answer that G has a valid 2-subcoloring; otherwise, we answer that G does not have a valid 2-subcoloring.

The algorithm is correct, because by the correctness of the rules there is an exhausted leaf of the spanning tree for which the associated constraint can be fulfilled if and only if G has a valid 2-subcoloring. Lemma 8.1 implies that an exhausted leaf will be accepted if and only if G has a valid 2-subcoloring.

We now analyze the running time. We first claim that the algorithm never executes Branching Rule 8.4 more than once in a root-leaf path of the search tree. By the description of the algorithm, none of the other branching rules are applicable if Branching Rule 8.4 is applied. Furthermore, it is clear from the description of the branching rules that none of the branching rules are applicable after Branching Rule 8.4 has been applied. Indeed, after Branching Rule 8.4 has been applied, observe that every vertex of $R^{\mathcal{C}}$ is in $N_{A_i^{\mathcal{C}}, B_j^{\mathcal{C}}}$ for some i, j for the associated constraint $\mathcal{C}$ which contradicts the prerequisites of the other branching rules. As for the reduction rules, clearly, neither Reduction Rule 8.1 nor Reduction Rule 8.2 can become applicable after applying Branching Rule 8.4. The claim follows, that is, Branching Rule 8.4 is applied at most once in a root-leaf path of the search tree.

To see that the search tree has $\mathcal{O}(4^k)$ nodes, note that Branching Rules 8.1, 8.2, and 8.3 are two-way branches, and each of these rules opens exactly one new cluster in each branch. By Reduction Rule 8.1, at most $k + 1$ clusters can be opened, and thus these branching rules lead to $\mathcal{O}(2^k)$ nodes of the search tree. After all these rules have been applied, possibly Branching Rule 8.4 will

be applied. Since Reduction Rule 8.1 cannot be applied, Branching Rule 8.4 yields at most 2^k branches. As shown above, no further branching rules will be applied after Branching Rule 8.4 has been applied. Hence, the search tree has $\mathcal{O}(4^k)$ leaves.

Finally, we analyze the running time along a root-leaf path in the search tree. We look at all rules individually and combine their running time bounds with Lemma 8.1. For Reduction Rule 8.1, clearly, whenever we open a cluster, we can check in $\mathcal{O}(1)$ time whether it becomes applicable and carry out the reduction rule accordingly. Hence, the time needed for Reduction Rule 8.1 is clearly bounded by $\mathcal{O}(n^2)$.

For Reduction Rule 8.2, note that it only becomes applicable to a vertex v if either a neighbor is put into one of the clusters in $A_*^{\mathcal{C}}$ or $B_*^{\mathcal{C}}$, or if a new cluster is opened with a neighbor of v. We claim that, whenever we open a cluster with a vertex u and whenever we put a vertex u into a cluster, we can decide in $\mathcal{O}(k \deg(u))$ time whether Reduction Rule 8.2 becomes applicable. The overall running time for checking and applying Reduction Rule 8.2 is then $\mathcal{O}(k(n+m))$ as each vertex is put into a cluster at most once. To decide whether Reduction Rule 8.2 becomes applicable when adding a vertex u to a cluster (or opening a cluster with u), we check for each neighbor w of u, whether it becomes adjacent to two clusters of $A_*^{\mathcal{C}}$ or of $B_*^{\mathcal{C}}$ or adjacent to a proper subset of the vertices of a cluster. This can be checked in $\mathcal{O}(k)$ time by maintaining, throughout the algorithm, the sizes of each cluster and, for each vertex $x \in R^{\mathcal{C}}$, two auxiliary arrays that keep track of which clusters x is adjacent to, and how many neighbors x has in each cluster. The overall update cost for these auxiliary arrays is $\mathcal{O}(n+m)$ because whenever we put a vertex into a cluster we can update the auxiliary arrays for all neighbors. Hence, overall we need $\mathcal{O}(k(n+m))$ time to check for applicability and for applying Reduction Rule 8.2, along a root-leaf path in the search tree.

Branching Rule 8.1 can clearly be applied in $\mathcal{O}(n+m)$ time and is applied at most k times, yielding $\mathcal{O}(k(n+m))$ time along a root-leaf path in the search tree.

Next, in the course of the algorithm we need to compute (and possibly recompute several times) the partition of $R^{\mathcal{C}}$ into the sets $N_{A_i^{\mathcal{C}}}$, $N_{B_j^{\mathcal{C}}}$ and $N_{A_i^{\mathcal{C}} B_j^{\mathcal{C}}}$. Using the information about the adjacency of vertices and clusters from the auxiliary arrays above, the partition can be computed in $\mathcal{O}(k(n+m))$ time. We recompute this partition whenever we have exhaustively applied Reduction Rule 8.1, Reduction Rule 8.2 and Branching Rule 8.1 and before we apply any of Branching Rule 8.2, Branching Rule 8.3 or Branching Rule 8.4. Since we showed above that the branching rules are applied at most $2k+1$ times, the time needed for computing the sets $N_{A_i^{\mathcal{C}}}$, $N_{B_j^{\mathcal{C}}}$ and $N_{A_i^{\mathcal{C}} B_j^{\mathcal{C}}}$ is $\mathcal{O}(k^2(n+m))$ along a root-leaf path in the search tree.

It is not hard to see that each of Branching Rule 8.2, Branching Rule 8.3, and Branching Rule 8.4 can be applied in $\mathcal{O}(n+m)$ time. As was shown above, they are applied at most $k+1$ times, yielding an $\mathcal{O}(k(n+m))$ time needed for these rules along a root-leaf path in the search tree. To conclude, we need

overall $\mathcal{O}(k^2(n+m))$ time to apply the reduction and branching rules along a root-leaf path in the search tree. Combining this with Lemma 8.1 and the fact that the search tree has $\mathcal{O}(4^k)$ leaves, we obtain an overall running time of $\mathcal{O}(4^k \cdot (n^2 + k^2(n+m))) = \mathcal{O}(4^k \cdot k^2 \cdot n^2)$. □

9 Hardness results

In this section, we present hardness results showing that significant improvements over the algorithms presented in the previous sections for Monopolar Recognition, 2-Subcoloring, and the (Π_A, Π_B)-Recognition problem in general, are unlikely. These hardness results are proved under the assumption P $\neq$ NP, or the stronger assumption that the Exponential Time Hypothesis (ETH) does not fail. We start with the following propositions showing that the existence of subexponential-time fixed-parameter algorithms for certain (Π_A, Π_B)-Recognition problems is unlikely:

Proposition 9.1. Let Π_A and Π_B be hereditary graph properties that can be characterized by a set of connected forbidden induced subgraphs and additionally Π_A is not the set of all edgeless graphs. Then (Π_A, Π_B)-Recognition cannot be solved in $2^{o(n+m)}$ time, unless the ETH fails.

Proof. For every Π_A and Π_B fulfilling the conditions in the proposition, Farrugia [12] presents a polynomial-time reduction from a variant of p-in-r SAT. Herein, we are given a boolean formula Φ with clauses of size r containing only positive literals, and the question is whether there is a truth assignment that sets exactly p variables in each clause to true. Given a formula Φ, the reduction of Farrugia [12] constructs an instance $G = (V, E)$ of (Π_A, Π_B)-Recognition where $n + m = \mathcal{O}(|\Phi|)$ (for each clause, the construction adds a gadget of constant size). The result now follows from the fact that p-in-r SAT cannot be solved in $2^{o(|\Phi|)}$ time, unless the ETH fails [17]. □

The lower bounds on the running time for Monopolar Recognition and 2-Subcoloring now follow from the above proposition:

Proposition 9.2. Monopolar Recognition parameterized by the number k of clusters in $G[A]$ and 2-Subcoloring parameterized by the total number k of clusters in $G[A]$ and $G[B]$ cannot be solved in $2^{o(k)}n^{\mathcal{O}(1)}$ time, unless the ETH fails.

Proof. Monopolar Recognition is the case of (Π_A, Π_B)-Recognition where Π_A defines the class of P_3-free graphs and Π_B defines the class of edgeless graphs. 2-Subcoloring is the case of (Π_A, Π_B)-Recognition where Π_A and Π_B define the class of P_3-free graphs. Hence, for both problems, Π_A and Π_B satisfy the conditions in Proposition 9.1. It now remains to observe that n is an upper bound on k in both cases. □

In Theorems 1.1 and 1.3 we give fixed-parameter algorithms for two (Π_A, Π_B)-Recognition problems, in both of which Π_A defines the set of all cluster

graphs, parameterized by the number of clusters in $G[A]$. Hence, one might hope for a generic fixed-parameter algorithm for such problems, irrespective of Π_B. However, polar graphs stand in our way. A graph $G = (V, E)$ has a *polar partition* if V can be partitioned into sets A and B such that $G[A]$ is a cluster graph and $G[B]$ is the complement of a cluster graph (a *co-cluster graph*) [27]. We have the following proposition:

Proposition 9.3. It is NP-hard to decide whether G has a polar partition (A, B) such that $G[A]$ is a cluster graph with one cluster or $G[B]$ is a co-cluster graph with one co-cluster.

Proof. Observe that a polar partition (A, B) of a graph $G = (V, E)$ where $G[B]$ is a co-cluster graph with a single co-cluster implies that $G[A]$ is a cluster graph and $G[B]$ is edgeless. Hence, G is monopolar. Since Monopolar Partition is NP-hard [12], it follows immediately that it is NP-hard to decide whether G has a polar partition (A, B) such that $G[B]$ is a co-cluster graph with a single co-cluster. Now observe that the complement of a polar graph is again a polar graph. Hence, a straightforward NP-hardness reduction by taking the complement reveals that it NP-hard to decide whether G has a polar partition (A, B) such that $G[A]$ is a cluster graph with a single cluster. □

10 Conclusion

In this paper, we developed the inductive recognition technique for designing algorithms for (Π_A, Π_B)-Recognition (graph) problems. Among other applications, we showed how this technique can be employed to design FPT algorithms for two graph partitioning problems: Monopolar Recognition parameterized by the number of cliques (in the cluster-graph part), and 2-Subcoloring parameterized by the smaller number of cliques between the two parts. These results generalize the well-known linear-time algorithm for recognizing split graphs and the polynomial-time algorithm for recognizing unipolar graphs, respectively. We believe that inductive recognition can be of general use for proving the fixed-parameter tractability of recognition problems that may not be amenable to other standard approaches in parameterized algorithmics. We also explored the boundaries of tractability for (Π_A, Π_B)-Recognition (graph) problems.

There are several open questions that ensue from our work. A natural concrete question is whether we can improve the running time of the FPTalgorithm for 2-Subcoloring with respect to k, the number of clusters $G[A]$. In particular, can we solve the problem in time $2^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$, or in time $f(k)n^2$? Note that the latter running time would match the quadratic running time for recognizing unipolar graphs, corresponding to the parameter value $k = 1$. It is also interesting to investigate further if we can obtain meta FPT-results for (Π_A, Π_B)-Recognition based on certain graph properties Π_A and Π_B, similar to the results obtained for mutually d-exclusive graph properties. Moreover, investigating the existence of polynomial kernels for Monopolar Recognition and 2-Subcoloring with respect to the considered parameters, and for

(Π_A, Π_B)-Recognition in general based on Π_A and Π_B and with respect to proper parameterizations, is certainly worth investigating.

Finally, one could consider vertex-partition problems that are defined using more than two graph properties, that is, $(\Pi_A, \Pi_B, \Pi_C, \ldots)$-Recognition. Observe that this problem is equivalent to $(\Pi_{A'}, \Pi_{B'})$-Recognition, where $\Pi_{A'} = \Pi_A$ and $\Pi_{B'}$ is the graph property "the vertices can be partitioned into sets that induce graphs with property Π_B, Π_C, ..., respectively". For example, a graph is d-subcolorable if its vertices can be partitioned into d sets that each induce a cluster graph, or alternatively, if its vertices can be partitioned into two sets, one of which is a cluster graph and the other of which is $(d-1)$-subcolorable. Using the aforementioned equivalence, the NP-hardness result of Farrugia [12] carries over to $(\Pi_A, \Pi_B, \Pi_C, \ldots)$-Recognition whenever each graph property is additive and hereditary.

The generalization of 2-Subcoloring to d-Subcoloring leaves the problem NP-hard per the above discussion; in fact, d-Subcoloring is NP-hard for any fixed $d \geq 2$ (this was explicitly proved by Achlioptas [1]). A stronger result follows from the observation that the straightforward reduction from d-Clique Cover (or d-Coloring) implies that d-Subcoloring is NP-hard for any fixed $d \geq 3$, even when the total number of clusters is at most d. This result implies that, unless P = NP, d-Subcoloring has no FPT or XP algorithm for any fixed $d \geq 3$ when parameterized by the total number of clusters. Hence, unless P = NP, Theorems 1.3 and 1.4 cannot be extended to d-Subcoloring for any fixed $d \geq 3$.

To extend the results for Monopolar Recognition and 2-Subcoloring, one can consider two different problems, both of which are NP-hard per the above discussion:

- Given a graph $G = (V, E)$ and $k \in \mathbb{N}$, can V be partitioned into a cluster graph with at most k clusters and two edgeless graphs?

  For $k = 1$, this problem asks whether there is an independent set in G whose removal leaves a split graph, and it can be solved in polynomial time [2]. Moreover, if Π_A is the property of being a cluster graph with at most k clusters and Π_B is the property of being a bipartite graph, then this problem is a (Π_A, Π_B)-Recognition problem. Since properties Π_A and Π_B are mutually $(2k+1)$-exclusive, Theorem 1.2 implies an XP algorithm for this problem when parameterized by k.

- Given a graph $G = (V, E)$ and $k \in \mathbb{N}$, can V be partitioned into two cluster graphs with at most k clusters in total and one edgeless graph?

  For $k = 1$, this problem asks whether G is a split graph, and it can be solved in polynomial time [15]. For $k \leq 2$, this problem asks whether there is an independent set in G whose removal leaves a co-bipartite graph, and it can be solved in polynomial time as well [2]. Moreover, if Π_A is the property of being 2-subcolorable with at most k clusters in total, and Π_B is the property of being an edgeless graph, then this problem is a (Π_A, Π_B)-Recognition problem. Since properties Π_A and Π_B are

mutually $(k+1)$-exclusive, Theorem 1.2 implies an XP algorithm for this problem when parameterized by k.

Since both problems have XP algorithms when parameterized by k, it is interesting to investigate whether any of them has an FPT algorithm. We leave these as open questions for future work.

References

[1] D. Achlioptas. The complexity of G-free colourability. *Discrete Math.*, 165–166:21–30, 1997.

[2] A. Brandstädt, V. B. Le, and T. Szymczak. The complexity of some problems related to GRAPH 3-COLORABILITY. *Discrete Appl. Math.*, 89:59–73, 1998.

[3] H. Broersma, F. V. Fomin, J. Nešetřil, and G. J. Woeginger. More about subcolorings. *Computing*, 69(3):187–203, 2002.

[4] S. Bruckner, F. Hüffner, and C. Komusiewicz. A graph modification approach for finding core-periphery structures in protein interaction networks. *Algorithms Mol. Biol.*, 10:16, 2015.

[5] R. Churchley and J. Huang. List monopolar partitions of claw-free graphs. *Discrete Math.*, 312(17):2545–2549, 2012.

[6] R. Churchley and J. Huang. On the polarity and monopolarity of graphs. *Journal of Graph Theory*, 76(2):138–148, 2014.

[7] M. Cygan, F. V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.

[8] R. Diestel. *Graph Theory, 4th Edition*. Springer, 2012.

[9] R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, Berlin, Heidelberg, 2013.

[10] T. Ekim, P. Hell, J. Stacho, and D. de Werra. Polarity of chordal graphs. *Discrete Appl. Math.*, 156(13):2469–2479, 2008.

[11] E. M. Eschen and X. Wang. Algorithms for unipolar and generalized split graphs. *Discrete Appl. Math.*, 162:195–201, 2014.

[12] A. Farrugia. Vertex-partitioning into fixed additive induced-hereditary properties is NP-hard. *Electron. J. Comb.*, 11(1):R46, 2004.

[13] J. Fiala, K. Jansen, V. B. Le, and E. Seidel. Graph subcolorings: Complexity and algorithms. *SIAM J. Discrete Math.*, 16(4):635–650, 2003.

[14] J. Gimbel and C. Hartman. Subcolorings and the subchromatic number of a graph. *Discrete Math.*, 272:139–154, 2003.

[15] P. L. Hammer and B. Simeone. The splittance of a graph. *Combinatorica*, 1(3):275–284, 1981.

[16] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.

[17] P. Jonsson, V. Lagerkvist, G. Nordh, and B. Zanuttini. Complexity of SAT problems, clone theory and the exponential time hypothesis. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '13)*, pages 1264–1277. SIAM, 2013.

[18] S. Kolay and F. Panolan. Parameterized Algorithms for Deletion to (r, ℓ)-Graphs. In *Proceedings of the 35th IARCS Annual Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS '15)*, volume 45 of *LIPIcs*, pages 420–433. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2015.

[19] J. Kratochvíl and I. Schiermeyer. On the computational complexity of $(\mathcal{O}, \mathcal{P})$-partition problems. *Discuss. Math. Graph Theory*, 17(2):253–258, 1997.

[20] V. B. Le and R. Nevries. Complexity and algorithms for recognizing polar and monopolar graphs. *Theor. Comput. Sci.*, 528:1–11, 2014.

[21] D. Lokshtanov. Wheel-free deletion is W[2]-hard. In *Proceedings of the Third International Workshop on Parameterized and Exact Computation (IWPEC '08)*, volume 5018 of *LNCS*, pages 141–147. Springer, 2008.

[22] C. McDiarmid and N. Yolov. Recognition of unipolar and generalised split graphs. *Algorithms*, 8(1):46–59, 2015.

[23] C. H. Papadimitriou. *Computational Complexity*. Addison Wesley, 1994.

[24] F. Ramsey. On a problem in formal logic. *London Mathematical Society*, 30(3):264–286, 1930.

[25] B. A. Reed, K. Smith, and A. Vetta. Finding odd cycle transversals. *Oper. Res. Lett.*, 32(4):299–301, 2004.

[26] J. Stacho. On 2-subcolourings of chordal graphs. In *Proceedings of the 8th Latin American Symposium on Theoretical Informatics (LATIN '08)*, volume 4957 of *LNCS*, pages 544–554. Springer, 2008.

[27] R. I. Tyshkevich and A. A. Chernyak. Decomposition of graphs. *Cybern. Syst. Anal.*, 21(2):231–242, 1985.

[28] R. I. Tyshkevich and A. A. Chernyak. Algorithms for the canonical decomposition of a graph and recognizing polarity. *Izvestia Akad. Nauk BSSR, ser. Fiz. Mat. Nauk*, 6:16–23, 1985. In Russian.